


NASA/CR-1998-208450



Semiannual Report

October 1, 1997 through March 31, 1998

*Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA*

Operated by Universities Space Research Association



National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Contract NAS1-97046

July 1998

Available from the following:

NASA Center for Aerospace Information (CASI)
7121 Standard Drive
Hanover, MD 21076-1320
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 487-4650

CONTENTS

	Page
Introduction	ii
Research in Progress	
Applied and Numerical Mathematics	1
Physical Sciences, Fluid Mechanics	7
Computer Science	14
Reports and Abstracts	21
ICASE Colloquia	36
ICASE Staff	38

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE)* is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U. S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis and algorithm development, fluid mechanics, and computer science in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Applied and numerical mathematics, including multidisciplinary design optimization;
- Theoretical and computational research in fluid mechanics in selected areas of interest to LaRC, such as transition, turbulence, flow control, and acoustics;
- Computer science: system software, systems engineering, and parallel algorithms.

ICASE reports are considered to be primarily preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period October 1, 1997 through March 31, 1998 is given in the Reports and Abstracts section which follows a brief description of the research in progress.

*ICASE is operated at NASA Langley Research Center, Hampton, VA, under the National Aeronautics and Space Administration, NASA Contract No. NAS1-97046. Financial support was provided by NASA Contract Nos. NAS1-97046, NAS1-19480, NAS1-18605, NAS1-18107, NAS1-17070, NAS1-17130, NAS1-15810, NAS1-16394, NAS1-14101, and NAS1-14472.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

BRIAN ALLAN

Reduced Order Model for Sensor/Actuator Problem

Placement of micro actuators and sensors for active flow control problems is complicated by the distributed nature of the flow equations and the dynamics of the closed-loop system. The goal of this research is to provide insight into the areas of the flow where actuation and sensing are important. This is done by performing an optimal feedback control design assuming distributed actuation and sensing everywhere in the flow field. The resulting optimal feedback kernel will then provide insight into the areas of the flow where actuation and sensing should be performed. Since the calculation of this optimal feedback kernel has a large computational cost, (order N^3 where N is the number of state variables) a reduced order model must be developed.

The flow is modeled by the unsteady incompressible Navier-Stokes equations in vorticity-stream function form. These equations are then linearized about a desired steady-state flow field and discretized using a finite difference scheme. A reduced order model is then generated from the linearized flow equations by constructing a Krylov subspace using the Arnoldi process. An implicitly restarted Arnoldi method is used in order to extract the leading eigenvectors from the Krylov subspace while avoiding storage problems. This Krylov subspace produces a reduced order model of the linearized flow equations used in the optimal feedback control design. We have successfully applied this method for a two-dimensional shear flow problem at low Reynolds numbers. This reduced order model is then used to calculate optimal feedback gains using an LQR approach. The resulting optimal feedback gains confirm that sensing and actuation in the shear region is important.

We are currently working on extending this research for the flow around an airfoil.

This work was performed in collaboration with J. Loncaric (ICASE).

EYAL ARIAN

Derivation of the Adjoint Equations for Incomplete Cost Functionals

The purpose of this project was to develop a general method of formulating proper boundary value problems from cost functionals considered “inadmissible” in the literature. We call such cost functionals “incomplete”.

We claim that for incomplete cost functionals additional auxiliary boundary equations are needed in the Lagrangian. These relations are obtained from the restriction of the interior PDE and its derivatives to the boundary. With these additional relations, proper cancellation of terms in the variation of the Lagrange functional can be obtained for any well defined cost functional. We demonstrate the derivation method on three problems involving the incomplete cost functionals using the potential, the compressible Euler, and the compressible NS equations.

This work has been completed and published in the ICASE Report No. 97-69, titled "Admitting the Inadmissible: Adjoint Formulation for Incomplete Cost Functionals in Aerodynamic Optimization."

This work was performed in collaboration with M.D. Salas (ICASE).

Development of Efficient Algorithms for Large Scale Aerodynamic Optimization

The purpose of this work is to develop algorithms which do not require more than a few full solutions of the flow equations to obtain the optimum.

We have completed the implementation of a new method on optimal shape design of airfoils, governed by the Euler equations, using pointwise discretization of the shape. The new method is based on preconditioning the gradient direction with an approximation of the inverse of the Hessian in the continuum level. This work has been published in the ICASE Report No. 98-14, titled "A Preconditioning Method for Shape Optimization Governed by the Euler Equations".

We intend to further investigate such methods in order to develop a robust preconditioning methods for large scale aerodynamic optimization.

This work was performed in collaboration with V. Vatsa (NASA Langley).

R. MICHAEL LEWIS

Pattern Search Methods for Nonlinear Optimization

Pattern search methods for nonlinear optimization have a number of features that make them attractive for use in engineering optimization. These methods are easy to understand and implement, are scalably parallel, and neither require nor estimate derivatives. However, heretofore provably robust pattern search methods did not exist for generally constrained optimization, a serious limitation in the context of engineering design.

We have undertaken the extension of pattern search methods to constrained optimization. We have completed the extension to linearly constrained problems via a feasible iterates approach. This we can do because the geometry of the boundary of the feasible region is completely known once the constraints are given. We have also extended pattern search methods to general nonlinear constraints via an augmented Lagrangian approach, adapting the algorithm of Conn, Gould, and Toint used in their LANCELOT code. This is a natural adaptation since this augmented Lagrangian method involves the solution of successive bound constrained minimization problems, and we have previously developed pattern search methods for the latter class of problems.

For purposes of comparison, we have also begun to investigate pattern search adaptations of both logarithmic and modified barrier methods.

This work was performed in collaboration with V. Torczon (The College of William & Mary) and the extension to barrier methods has been begun together with S. Nash (George Mason University).

A Posteriori Finite Element Bounds for Sensitivity Calculations

In the optimization of systems governed by differential equations one would like to use the coarsest mesh possible at any given step so as to reduce the cost of the optimization iteration. In a recent series of papers, Patera, Peraire, and their collaborators have presented an *a posteriori* approach to computing quantitative bounds on the mesh dependence of certain functionals of the solutions of differential equations. We have begun to apply these ideas in the context of optimization.

We have begun with developing *a posteriori* bounds for sensitivities of output linear functionals with respect to various parameters (such as coefficients) in boundary-value problems. Using either the sensitivity equations or adjoint equations one can write the output's sensitivity as a functional of the solution of a system of differential equations. Using the general theory, one can then derive suitable bounds on the error in the sensitivities on a coarse grid relative to a finer grid. Numerical results indicate that the bounds can be quite good.

We are currently extending the general approach to broader class of output functionals and differential equations, as well as investigating using approximate function values with error bounds in optimization.

This work was performed in collaboration with T. Patera and J. Peraire (Massachusetts Institute of Technology).

Decomposition Approaches in MDO

A number of approaches to formulating multidisciplinary optimization problems have been proposed. These approaches attempt to address the problem of capturing the coupling of disciplines while preserving the autonomy with which they perform the computations necessary for the overall design optimization process. However, while there may be a mathematical equivalence between different formulations of MDO problems, some formulations cause difficulties for optimization algorithms. We have begun a systematic analysis of various formulations, beginning with Collaborative Optimization (CO) approach developed at Stanford.

CO can be viewed as a special sort of bilevel optimization problem. We have determined that CO has a number of analytical features that will cause problems for numerical optimization algorithms. For instance, some versions of CO lead to optimization problems for which *no* constraint qualification (e.g., calmness) holds, and no Lagrange multipliers exist for the system-level constraints. In addition there exist a number of practical problems with CO that derive from its nature as a bilevel optimization problem.

From our understanding of the strengths and weaknesses of Collaborative Optimization we hope to suggest improvements. We also intend to analyze other methods that have been proposed.

This work was performed in collaboration with N. Alexandrov (NASA Langley).

JOSIP LONCARIC

Spatial Structure of Optimal Flow Control Around Airfoil

Designing distributed control systems begins with the sensor/actuator placement problem. While in some situations discrete search of combinatorial complexity seems unavoidable, continuum problems suggest solving a related question. *If one could sense everything and actuate everywhere, what should one do?* The answer to this question has polynomial complexity (of order N^3 where N is the number of state variables) and can serve as the initial effectiveness filter capable of rejecting a large portion of the design search space. This favorable situation can have several causes. In our investigation we focus on the effect of no-slip boundary conditions on an optimal flow control of the unsteady Stokes flow around the NACA 0015 airfoil. This test problem aims at the development of numerical methods capable of solving the problem of stabilizing the desired flow around wings at low to moderate Reynolds numbers.

Our approach begins by defining a *pseudodifferential representation* ξ of the flow. This step is similar to the Wiener-Hopf factorization of the Laplacian. We have shown that for the flow around a cylinder this leads to an explicit diagonalization of the system dynamics $\xi_t = \mathcal{A}\xi$ by means of the Fourier and Weber transforms. We then pose and solve an optimal distributed LQR problem with gain ϵ . A rational approximation to the

optimal feedback kernel is derived and shown to perform within 0.026 percent of the exact optimum even in the worst case. Using the vorticity representation in conformally mapped geometries, this approximation is decomposed into the analytic free space solution and a boundary term which can be evaluated numerically.

A numerical procedure to display the spatial form of the nearly optimal feedback in response to a disturbance vortex is under development. While the local vorticity damping is a part of this solution, the boundary term part of the optimal actuation is of particular interest. The insight gained in this study will provide guidance for sensor/actuator placement as we extend the control theoretic approach to low and moderate Reynolds number flows.

LI-SHI LUO

Lattice Boltzmann Model for Non-Ideal Gases

The key issues in the study of multi-phase (*e.g.*, liquid-vapor) flows are the modeling of interfaces and phase transition among different phases. It is difficult to use the Navier-Stokes equations to model the inhomogeneous multi-phase flows because the interfacial tracking is a laborious computation. In the past few years, a number of lattice Boltzmann models have been developed to model multi-phase flows. However, the multi-phase lattice Boltzmann equation is still lacking a rigorous theoretical basis. For instance, previous multi-phase lattice Boltzmann models do *not* have a consistent equilibrium thermodynamics. The present work applies the Enskog Theory of hard spheres to revise the theory of the multi-phase lattice Boltzmann equation.

With the Enskog theory we were able to derive a new multi-phase lattice Boltzmann model which has a consistent equilibrium thermodynamics. We have rigorously demonstrated the deficiencies in the previous multi-phase lattice Boltzmann models and provided a systematic procedure to derive *correct* multi-phase lattice Boltzmann model based upon the Enskog theory (or the revised Enskog theory). The present work has been submitted to Physical Review Letters and an ICASE report is in preparation.

We intend to derive thermodynamically consistent multi-component lattice Boltzmann model in the future based upon the same methodology.

Lattice Boltzmann Method and Beam Scheme

The mathematical connection of the method of the lattice Boltzmann equation (LBE) and the beam scheme is being studied.

We carefully analyze the LBE method and the beam scheme, and derive their mathematical similarities, and compare their valid regime of applicability. This work has been submitted to Physical Review E and an ICASE report has also been submitted.

We plan to develop a finite-volume lattice Boltzmann scheme based upon the insights gained in this work.

This work was performed in collaboration with K. Xu (Hong Kong Science and Technology University).

DIMITRI MAVRIPLIS

Three-Dimensional Unstructured Multigrid Convergence Acceleration for Highly Stretched Meshes

For high-Reynolds number viscous flow simulations, efficient resolution of the thin boundary layer and wake regions requires mesh spacings several orders of magnitude smaller in the normal direction than in the

streamwise direction. This extreme grid stretching results in poor multigrid convergence rates, usually one to two orders of magnitude slower than those observed for equivalent inviscid flow problems without grid stretching. The purpose of this work is to devise improved multigrid techniques for acceleration convergence on anisotropic three dimensional grids.

The approach taken consists of using implicit line solvers in the direction normal to the grid stretching combined with semi-coarsening or directional coarsening multigrid to alleviate the stiffness due to grid anisotropy. A graph algorithm has been implemented which combines edges of the original mesh into lines which follow the direction of maximum coupling in the unstructured mesh. A similar algorithm is used to selectively coarsen the original fine mesh by removing points along the directions of strong coupling, thus recursively generating a hierarchical set of coarse meshes for the multigrid algorithm. These coarse levels are constructed using the method of agglomeration. While the original isotropic agglomeration method was based on an unweighted-graph algorithm, the new directional agglomeration procedure relies on a weighted-graph algorithm. Preconditioning techniques to alleviate low-Mach number induced stiffness in the equations have also been implemented. Finally, a Newton GMRES technique has been implemented which uses the above described multigrid and preconditioning techniques as a preconditioner themselves. Using this approach, convergence rates for viscous flows which are independent of the degree of mesh stretching can be obtained. Results have been obtained in three dimensions using grids up to 2 million points and demonstrating order of magnitude increase in asymptotic convergence rates over the original explicit isotropic agglomeration algorithm.

The three-dimensional multigrid solver is currently being ported to distributed memory parallel computer architecture. This will permit the simulation of high-lift flows on grids of up to 20 million points in the near future.

CHRISTINE TOOMER

Sensitivity-Based Approximation Methods for Use in Multidisciplinary Design Optimization Studies

Aerodynamic optimization studies require a database of good quality information relevant to the design space of a chosen geometry. The traditional route for obtaining this data is from multiple runs of a Computational Fluid Dynamics (CFD) code. However, such an approach (especially in viscous, turbulent mode) on full-aircraft multiblock geometries is still too expensive to be used within the design environment. Sensitivity-based approximation methods can significantly cut the time required to produce this data. This investigation continues the development of a direct sensitivity code which uses what has been termed a Projected Implicit Reconstruction (PIR) algorithm. This type of model reconstructs the flowfield around a chosen geometry using information from the previous design point. As the code sweeps through a design space it stores information from each design point. The aim of this investigation is to enhance the code such that it produces good quality operating conditions and geometric data for viscous, turbulent, whole aircraft geometries considerably faster than a CFD code.

During the last three months, the accuracy of the inviscid model was considerably improved by implementing a second-order scheme into the flux gradient calculations. This is particularly effective for Mach number sensitivity studies where not only has the relative error decreased by two-thirds, but the run-times have reduced by over 25%. The inviscid model, when compared to using a CFD code for data generation, produces savings in time between 60% and 90%, depending on the case and the design parameter(s) considered. The viscous, laminar model is also under test and improvements have been made to the stress

algorithms used in the flux gradient calculation. A standard two-equation turbulence model is currently being added into the code.

The future work involves finalizing the viscous capability and amending the code to perform more extensive grid sensitivity studies. Currently only one response surface method has been used to map the discrete data results across the design space. Investigations into the suitability of other methods will be undertaken when optimizing typical aircraft geometries over flight envelopes.

PHYSICAL SCIENCES, FLUID MECHANICS

P. BALAKUMAR

Nonlinear and Secondary Stability of Three-Dimensional Boundary Layers

Despite numerous efforts towards a detailed understanding of the fundamental physics of the laminar-turbulent breakdown in three-dimensional incompressible boundary layers, the current picture of the transition process is not yet complete, and there is a lack of efficient and robust codes for the application to large-scale and industry-oriented problems, like the transition prediction on infinite swept wings. The objective of this research is to develop efficient and robust numerical methods that capture the flow physics of the present problem accurately, and to gain a deeper qualitative and quantitative insight into the instability mechanisms involved in the transition process in crossflow instability dominated boundary layers.

The current approach to this problem involves solving the nonlinear PSE explicitly for an investigation of the nonlinear stability, thereby creating input data for the temporal examination of the secondary instability in the region of nonlinear amplitude saturation applying Floquet theory. Using the quantitative information from the temporal secondary instability analysis, the secondary instability is then modeled in a more straightforward manner using the nonlinear PSE. The work performed thus far indicates that the taken approach is capable of capturing the nonlinear primary and secondary instability in one computational run using the PSE. In parallel to this approach, a novel implicit method for solving the nonlinear PSE has been developed. This method is much more robust and efficient in the region of nonlinear amplitude saturation. Now that the implicit PSE-solver has been validated, it will be applied to several problems for which the explicit solver does not converge.

For future research, it is planned to improve the efficiency of the implicit PSE-solver further, and apply it to modeling the complete path of the laminar turbulent transition in two- and three-dimensional boundary layer flows.

YA-CHIN CHEN

New Adaptive Technique for Signal Prediction with Applications to Speech Coding and Input Identification

Signal prediction and data compression are known to be extremely difficult, if not impossible, for systems without knowing the excitation signal. It requires both system modeling and input estimation simultaneously. The objective is to develop a signal prediction technique for applications such as speech coding and synthesis and input identification.

The approach is based on the disciplines of signal processing and system identification to estimate the input excitation. It uses the multi-modulus blind equalization method in conjunction with the geometric subspace projection theory to formulate the basic prediction equation. The coefficients of the blind equalizer and the input excitation are solved simultaneously and recursively by a conventional recursive least-squares algorithm. The excitation codebook associated with the excitation signal is generated analytically without the stochastic assumptions. The parameters to represent the signal include the coefficients of the blind equalizer and excitation and the index of the codebook.

The algorithm has been completed in MATLAB 4.2 and speech coding (analysis and synthesis) has been successfully performed. The proposed vocoder can perform well with an 8-bit codebook coding, whereas the conventional vocoders usually need 10-bit codebook coding to guarantee a good quality.

The new technique is a significant improvement over the conventional methods such as Linear Predictive Coding that is one of vocoders widely used in practice today. The algorithm is recursive and thus feasible for real-time implementation in practice. The use of geometry space concepts lead to a complete explanation of the proposed technique. In addition, a variable-frame concept is proposed to block the same excitation signal in sequence in order to further reduce the storage space and increase the transmission rate.

Other work of interest being examined includes excitation codebook generation. The technique will be applied to a number of control-related problems, such as disturbance identification.

This was performed in collaboration with J.-N. Juang (NASA Langley).

SANG-HYON CHU

Development of Microwave-Driven Smart Material Actuator

“Wireless” actuators controlled remotely with microwaves offer tremendous advantages over hard-wired actuators, especially for space applications such as the Next Generation Space Telescope (NGST), in which thousands of discrete actuators are required to effect high precision distributed shape-control of the primary reflector. This new concept alleviates the need for hard-wire connections resulting in significantly simpler system designs and lower system mass. “Wireless” control of an actuator involves constructing actuator elements with layers of metal conductors embedded into smart material. The H-field of microwave that is incident on a conductor instantaneously induces a current or an electric field between metal conductors. The desired motion or displacement of actuators is, therefore, made possible by controlling an electric potential due to the coupled microwave energy. In this work, a single microwave-driven actuator will be developed and tested for the proof-of-principle.

In order to demonstrate the application of microwaves as a power source to activate and control a smart material actuator, an anechoic chamber and a test specimen have been prepared for the test. The specimen is a pair of bare metal strips and a Schottky diode is used as a rectifier of AC field that is generated by microwave. This study as the first stage of a whole project is expected to play a very important role in understanding the fundamental issues of microwave-driven actuator concepts.

After mapping-out of electromagnetic field in an anechoic chamber is completed, the test specimen will be set up for the far-field exposure of a 10 GHz and 2-watt microwave source. The size, gap distance of metal strips, and wave incident angle will be varied to determine the optimum electric potential in case of various levels of microwave. Another specimen, a metal strips-embedded smart material actuator, will be tested under the same condition as a pair of bare metal strips. The test results of these two specimens will be compared to define the effects of smart material’s dielectric constant.

SHARATH GIRIMAJI

Non-Equilibrium Complex Deformation Turbulence: Pressure-Strain Modeling

The pressure-strain correlation model developed in the previous year was applied to a variety of complex deformation problems with success. The most notable accomplishment of this model is its ability to distinguish between rotating reference frame flows and flows with mean flow rotation. Previous models treated the two flows in an identical manner leading to a good prediction of rotating reference frame flows while the mean-flow rotation predictions were completely, even qualitatively, wrong. The present model, on the other hand, yields kinetic energy decay for rotating homogeneous flows and energy growth for elliptic flows

as required. This enhanced capability of the model is made possible by forcing consistency with a recently derived theorem on the behavior of velocity fluctuations in general rotating flows.

Velocity Fluctuations in Rapidly Rotating Turbulence

Many flows of engineering, especially aerodynamic, interest are dominated by the effects of mean-flow rotation. However, the behavior of velocity fluctuations in all flows with nearly identical mean flows can be vastly different. The exact nature of the velocity fluctuations depends upon the manner in which the fluctuations are advected by the mean flow. This important issue was identified and quantified for the first time. The manner of advection of the velocity fluctuations may be what distinguishes the reference-frame rotation flows (occurring in geophysical problems) from the mean-flow rotation flows encountered in the wake of aircrafts (e.g., wing-tip vortex). A theorem that relates the pressure fluctuations to the velocity fluctuations in a general rotating flow was derived. This theorem can be viewed as a generalization of the Taylor-Proudman theorem which is applicable only for reference-frame rotation flows.

This theorem has already been employed in deriving the first pressure-strain correlation models which can distinguish between various types of rotating flows.

This work was performed in collaboration with J.R. Ristorcelli (ICASE).

Non-Equilibrium Turbulence: Algebraic Reynolds Stress Modeling

The new algebraic approximation of the differential derivative, postulated in the previous research period, was rederived with more mathematical rigor. This approximation was also subject to extensive testing and validation in a variety of homogeneous turbulent flows. With the exception of early transient behavior in rapidly distorted flows, the model approximation appears to be quite reasonable. In the early transient region, it appears that another variable, apart from kinetic energy and dissipation is required to quantify the degree of departure from the slow-manifold stage of turbulence. These issues are being currently considered.

CHESTER GROSCH

Simulation of Supersonic Jet Mixing

Mixing enhancement of high and low speed streams is utilized as a means to improve efficiency of supersonic combustors, reduce aircraft signatures, and control high speed jet noise. One common method of mixing enhancement has been the introduction of streamwise vorticity by prism shaped wedges. Such devices have been investigated experimentally but no theoretical or semi-empirical method exists that can relate tab design to enhanced mixing rate and noise reduction. Another common method of mixing enhancement is to use lobe mixer ejectors. Again, only experimental information is available to evaluate the performance and guide the design of these mixers. The objective of this research is to use numerical simulation to examine the performance of both these devices in order to understand the physics of the mixing and how it is affected by changes in the parameters of these devices.

A set of numerical calculations were carried out using the compressible three-dimensional Navier-Stokes equations. The tabs were modeled by pairs of counter rotating vortices. Various geometric configurations of the lobe mixers are being used together with periodic side boundary conditions to simulate an array of these devices.

Simulations with tabs used an axisymmetric round inflow for the jet with no tabs, two tabs and four tabs arranged symmetrically on the edge of the jet. The results of the simulations show a striking similarity

in the flow evolution between numerical and experimental jets, particularly in the jet potential core region. The experimental and computational results also show similar values for mass entrainment. Because of complications associated with computational boundaries, the comparison between experiment and simulation degrades beyond 8 jet diameters downstream. These results will be presented by my co-author, J.M. Seiner, at the 4th AIAA/CEAS Aeroacoustics Conference in Toulouse, France, 2-4 June 1998.

The initial simulation of the lobe mixer flow shows that the jet becomes unstable and oscillates in the “garden hose” mode. The frequency of this mode is constant throughout the entire simulation (24,000 time steps) and equal to 16 KHz with a Strouhal number of about 0.44. The disturbances grow with downstream distance. The no slip boundary condition at the top and bottom walls cause large shear which generates streamwise vorticity. As the disturbances become nonlinear the streamwise vorticity migrates from the solid boundaries into the interior. Strong mixing follows and, by about halfway down the channel, the jet and coflow become nearly fully mixed.

Future experimental and numerical studies are required to more clearly define the initial induced vorticity field in the round jet. Future experiments will use PIV imaging to measure the cross-stream vectors. The experimental data will be used to set the inflow conditions for the numerical simulations. In addition new simulations will be performed in a larger computational domain in order to avoid boundary interference noted in the present study at large downstream distances.

Further calculations are planned for the lobe mixer including varying the geometry and adding tabs on the sides of the lobes.

ALEX POVITSKY

Computation of Three-Dimensional Acoustic Fields

Methods of direct solving of large bandwidth linear systems are widely used in computational mechanics. Parallelization of these solvers is hindered by global spatial data dependencies leading to pipelining of parallel algorithms. The main disadvantage is that during the pipeline process there are a lot of communications between processors at each time step and some processors will be idle at the switch from forward to backward steps of the Thomas algorithm and at the end of a time step. Implicit CFD and compact CAA solvers, including those involved in this project, use directional splitting methods, which are based on inversion of band matrices in three directions. In this investigation, we will develop a parallel version of a CFD-CAA solver for computations of noise from an aircraft providing exactly the same solution as the original serial solver and keeping low parallelization penalty.

The following question has been addressed at the beginning of this research. Is it possible to improve the parallelization efficiency of the Thomas algorithm? To answer this question, the two-step (forward-backward) pipeline algorithms (PA) have been introduced formally. For any PA algorithm, the minimum idle time of the i^{th} processor is equal to $2(P - i)$, where P is the number of processors. The similar results are valid for extended PA, where lines are gathered in different portions for forward and backward computations. Thus, it is impossible to reduce idle time of processors involved in PA computations. To cure the problem of processors' idle time, a new parallel version of the Thomas algorithm has been proposed and denoted as the Immediate Backward Thomas Pipeline Algorithm (IB-TPA). This formulation of the Thomas algorithm has completed data for next computations before idle of processor and may realize next computations while processors are idle. In turn, the IB-TPA can solve more lines per one message and reduce overall communication latency time. To make the IB-TPA feasible, algorithms for assignment of

processors computations and communications schedule have been derived. The IB-TPA has been developed and verified in modern MIMD computers, SP2 and CRAY T3E, for a test problem using a three-dimensional ADI method for a partial differential equation. The obtained parallelization penalty is in good agreement with our theoretical predictions for one-dimensional multi-domain partitioning. Extension of the IB-TPA for two- and three- dimensional partitioning is in progress. Now that the basic formulation of parallel Thomas algorithm and its use in ADI methods are understood, we can implement it to serial CFD-CAA codes.

A typical large-scale code includes dozens of routines, which are executed prior to the band matrix solver. Hand parallelization of these routines usually takes months of full-time work of a professional programmer or scientist. The use of automated parallelization tools can automate approximately 90% of the labor. In order to find a best tool for semi-automatic parallelization, the Computer Aided Parallelization Tools (CAPTools) and High Performance Fortran (HPF) has been studied and examined. The main advantage of the CAPTools over the HPF is that the CAPTools can track inter-routine data dependence and avoid unnecessary inter-processor streams of data. In collaboration with Parallel Processing Group at University of Greenwich (CAPTools developer) we will implement and update CAPTools for parallelization of the FDL3DI code, developed by the Wright-Patterson Air Force Base (WPAFB), and the 6th order compact CAA code, developed by Dr. M. Visbal, technical area leader, WPAFB.

J.R. RISTORCELLI

Compressible Turbulence: The Pressure-Strain Covariance

Models for the effects of the compressibility of turbulence are of crucial importance for the prediction of turbulent compressible flows.

A systematic perturbation procedure has been developed and used to develop models for the accumulated effects of compressibility on the moment evolution equations of compressible turbulence consistent with the present levels of turbulence models in use ($k - \epsilon$). Models for the pressure-dilatation ($\langle pd \rangle$) derived in earlier work have been used to develop models for the compressible pressure-strain covariance. This is possible in the context of a linear pressure-strain model. The fact that the pressure-dilatation is the first invariant of the pressure-strain has been used. The compressible pressure-strain model, which contains no empirical calibration constants, has been used to calculate the mixing layer and the compressible jet flows.

Theory and computations have shown that the effects of the pressure-dilatation are negligible for both the mixing layer and the jet flows. However, a compressible pressure-strain model using the asymptotic model for the pressure-dilatation does, in a manner consistent with experimental data, reduce both the Reynolds stresses and the production of the turbulence energy.

This work was performed in collaboration with V. Adumitroaie (SUNY Buffalo).

Aeroacoustics

Our objective is to devise models to predict the sound power radiated from the broadbanded fluctuations in turbulent flows using two-equation and higher order turbulence models.

This requires the use of the methods of statistical fluid mechanics in developing a predictive scheme which will have as inputs turbulence quantities calculated in single-point turbulence closures: the Reynolds stress anisotropy, the kinetic energy, and the spectral cascade rate [dissipation]. This procedure can be useful only if rational models for the two-point spatial and temporal correlations are devised. To this end, work has focused on kinematically appropriate models for temporal and spatial correlations.

The work is being extended to the difficult problem of anisotropic turbulence — the situation characteristic of the engineering problem.

ROBERT RUBINSTEIN

Rotating Compressible Flow

In compressible jet flows, the reduced mixing caused by compressibility can be undesirable. Swirl has been proposed as a means to counteract this effect; compressible swirling jets were the object of a large experimental program at NASA Langley. The analysis of such flows poses the problem of coupled effects of rotation and compressibility on turbulence. The occurrence of two-dimensionless parameters: the Mach number and the Rossby number, precludes the derivation of turbulence models by purely heuristic methods.

We analyzed the combined effects of rotation and compressibility on turbulence by combining our earlier work on compressibility with Zhou's phenomenological theory of rotating turbulence. The analysis applied Yoshizawa's two-scale formalism; but instead of perturbing about a Kolmogorov steady state, as is usual in these calculations, we perturb about the steady state for strongly rotating turbulence proposed by Zhou. Results are given both for turbulence transport models and for compressibility and rotation-modified subgrid models for large-eddy simulations. We find that certain compressibility effects, the enhanced sound speed and the effective turbulent pressure, are, in fact, reduced by rotation.

A natural extension of this work will be to examine the effect of rotation on compressible mixing. This effect does not occur at the order of perturbation theory analyzed here.

This work was performed in collaboration with Y. Zhou (ICASE and IBM) and G. Erlebacher (Florida State University).

Shock Wave Propagation in Weakly Ionized Gases

Recent proposals for a high-speed civil transport have revived interest in the experimentally observed mitigation of shock waves in the presence of charged particles. Since experiments suggest that the addition of even trace quantities of charged particles can modify shock properties significantly, this effect might provide an energy-efficient means to fly vehicles designed for subsonic flight at supersonic speeds.

Previous research has shown that elementary two-fluid models of this type of system greatly underpredict the observed effects. These models assume that the neutral and plasma components are separate in thermal equilibrium; our work starts from the more realistic assumption that the external energy source maintains this system in a steady state far from equilibrium. Since the distribution of molecular energy in such a state differs greatly from the energy distribution in thermal equilibrium, it is possible that the properties of shock and sound waves in such systems may be altered accordingly.

We find that a hard-sphere fluid admits a power-law distribution function when particles are injected at high energy and removed at low energy. We are currently simulating this steady state using a direct simulation Monte Carlo code due to G. Bird. If the code can reproduce this steady state, we will use it to study the properties of shock waves propagating through this fluid. Work on corresponding two-component models and on lattice Boltzmann representations of such fluids is also in progress.

This work was performed in collaboration with A.H. Auslender (NASA Langley).

Turbulent flows subjected to solid-body rotation occur in many engineering and geophysical applications, such as in turbomachinery with rotating blades, and in problems involving the rotation of the Earth. The effects of rotation (through a Coriolis force) on turbulence structure are known to be subtle, yet profound: although it has no direct role in the kinetic energy budget, rotation weakens the fundamental property of an energy cascade from the large scales to the small scales. Our objectives are to study the response of isotropic turbulence under uniform solid-body rotation in the presence of external energy input at the large scales, and to examine the spectral dynamics underlying this response. A substantial part of this work is devoted to a comprehensive description of the evolution of different anisotropy measures in both physical and spectral space. We attempt to explain the observed behavior through analyses of Reynolds stress budgets and triadic energy transfer, where differences with non-rotating turbulence will be emphasized. In particular, at both low and high wavenumbers, we distinguish between energy transfer for velocity components parallel and perpendicular to the axis of rotation.

We have used direct numerical simulations to study the response of initially isotropic turbulence subjected to uniform solid-body rotation in the presence of external energy input by numerical forcing. The grid resolution is 256^3 , initial Taylor scale Reynolds number is 140 (sufficient for a limited $k^{-5/3}$ inertial range in the energy spectrum), and rotation rates correspond to initial turbulent Rossby numbers of 0.0195 and 0.0039. A major emphasis is to examine the detailed structure and anisotropy characteristics of spectral transfer, as well as the possibility of deviations from local isotropy at the small scales. As expected, energy transfer and hence dissipation are greatly reduced by rotation. As predicted, the energy spectrum attains a characteristic k^{-2} scaling range at intermediate wavenumbers, although the Reynolds numbers in this work are presumably not high enough for an accurate determination of parameters characterizing a rotation-modified inertial range. Several measures of anisotropy are examined. Analysis of component budgets show that pressure-strain correlations and temporal oscillations due to rotation prevent the Reynolds stress tensor (which is a one-point characteristic) from developing significant anisotropy over time. However, strong anisotropy is evident in (i) increased integral length scales along the axis of rotation, (ii) the high wavenumber part of the radial energy spectrum, and (iii) deviations of one-dimensional longitudinal and transverse spectra from isotropy relations at high wavenumbers. In addition, vorticity statistics exhibit anisotropic properties suggestive of the formation of slender vortical structures aligned with the axis of rotation. Furthermore, spectral transfer properties are studied in some detail, mainly in terms of triadic interactions in wavenumber space. As expected, rotation greatly weakens the energy cascade to the small scales. Considerable differences are found between the transfer characteristics of the velocity components parallel and perpendicular to the axis of rotation. The former has similarities with results in non-rotating turbulence and contributes dominantly to the reduced energy cascade. By contrast, the latter shows indications of reverse energy transfer from the intermediate scales towards the lowest wavenumbers in the spectrum, primarily as a result of our finding that the triadic interaction scales are local.

We plan to use new understanding from this work to develop and refine engineering models for turbulent flows subject to strong rotations.

This work was performed in collaboration with P.K. Yeung (Georgia Tech).

COMPUTER SCIENCE

TOM CROCKETT

Portable Parallel Rendering Algorithms and Software for Scientific Visualization Applications

During the past year, NASA has phased out older systems and is installing a new generation of parallel supercomputers, including several 64- and 128-processor SGI/Cray Origin2000's, a 256-processor HP Exemplar, and a 1024-processor Cray T3E. Although these new systems continue to support message-passing as a programming paradigm, they also provide architectural support for shared memory and/or remote memory accesses. To use these systems effectively in visualization applications, we need to (1) determine how well existing parallel rendering algorithms and software will perform on the new architectures, (2) assess scalability with higher processor counts, (3) develop and test new algorithms if necessary, and (4) integrate algorithmic innovations into portable parallel rendering and visualization software which will run across a range of architectural platforms.

Over the last few months we have begun to study these issues by porting our PGL rendering system to two "distributed shared memory" systems, the Exemplar and Origin2000. We are conducting experiments with our existing MPI-based rendering algorithms to determine whether the elusive goal of "portable parallel performance" can be met without extensive platform-dependent recoding. Results on the Exemplar were generally quite good with up to 64 processors, but beyond that performance collapsed, yielding rendering rates for 128 processors which were little or no better than those observed on the five-year-old Intel Paragon. The precise cause is still under investigation, although our internal performance instrumentation indicates excessive message-passing overheads. We have provided a copy of our code and a detailed description of our results to HP for their analysis. We are currently repeating our experiments on the Origin2000 to see if a similar pattern is observed.

We have also been studying scalability issues, and have designed a two-step communication algorithm for use in PGL. This approach, first demonstrated by Ellsworth, trades off expensive message-passing for inexpensive data copying. While it reduces the number of messages generated, it also increases bandwidth demands on the communication network and on local memory. Thus testing on a variety of platforms will be essential in characterizing its overall effectiveness. Our current termination detection protocol proved to be incompatible with the two-step strategy, so we designed and implemented a new protocol which is conceptually simpler and much more scalable. In benchmark tests on the Intel Paragon, the new termination algorithm increased performance by 17% with 128 processors, and by 80% with 256 and 512 processors.

Our next step is to port PGL to the Cray T3E and assess performance on that platform. We then plan to implement and test the two-step algorithm on a variety of systems, and to compare explicit shared-memory algorithms against the current message-passing approach on distributed shared memory systems.

STEPHEN GUATTERY

Pseudoinverses of Symmetric Matrices and Generalized Graph Embeddings

Connections between Laplacian spectra (specifically the second smallest eigenvalue λ_2 and its corresponding eigenvector, or, in the case of a zero Dirichlet boundary, the smallest eigenvalue λ_1 and its eigenvector) and properties of the corresponding graphs have applications in algorithms. A number of graph embedding

techniques that give lower bounds on the smallest nontrivial eigenvalues of Laplacians exist (an example is the path resistance method recently introduced by Steve Guattery (ICASE), Tom Leighton (MIT), and Gary Miller (CMU), which works for both the graph (positive semidefinite) and zero boundary (positive definite) cases. A common attribute of embedding techniques is that they do not provide tight lower bounds; Guattery and Miller have recently shown that the gap in these bounds is a result of the representation of the problem. They showed that, by slightly modifying the representation, it is possible to construct a specific embedding, the current flow embedding, such that the matrix representation of the embedding can be used to construct a matrix whose eigenvalues and eigenvectors have an exact relationship to those of the Laplacian. In the Dirichlet boundary case, they showed that the embedding matrix can be used to construct the inverse of the Laplacian. These results have an interpretation in terms of resistive circuits and Kirchoff's and Ohm's laws.

We have now generalized this result to show that, for any real symmetric matrix, its pseudoinverse can be constructed in terms of a generalization of the current flow embedding. The resistive circuit analogy is lost in this case, though the embedding still obeys generalizations of the electrical laws.

In current work, we are considering whether approximations to the generalized current flow embedding can be used to construct useful preconditioners for such systems.

VICTORIA INTERRANTE

Feature Identification and Enhancement for 3D Flow Visualization

Research has focused on the development of perceptually-based methods for more effectively visualizing 3D datasets. The approach is based on looking for insight into the scientific principles that underlie the art of effective visual representation, based on fundamental findings in human visual perception and observation of practices in art and illustration.

Recent work has been directed at designing techniques for more effectively conveying 3D flow with line integral convolution texture. Accomplishments include:

- developing an efficient new method for enhancing depth discontinuities between overlapping streamlines in an arbitrary projection through the use of visibility-impeding halos,
- demonstrating effective uses of color parameters such as hue and saturation for conveying information about scalar quantities over a 3D flow,
- deriving an efficient technique for portraying directional information in a static volume-rendered representation of a 3D flow based on the use of exponentially defined filter kernels in the progressive integration step of LIC.

My future work plan includes extending the directional LIC technique to unsteady flow data, investigating techniques for more effectively conveying the relative 3D positions of scattered elements in a static flow, and refining methods for extracting shock surfaces from 3D data using a gradient magnitude ridge finding algorithm.

DAVID KEYES

Architecturally Adapted Algorithms for Parallel CFD

Newton-Krylov-Schwarz (NKS) methods permit the data locality natively present in most PDE-based applications (CFD, CSM, CEM, etc.) to be exploited on hierarchical distributed memory systems while delivering machine-zero satisfaction of the discrete conservation laws, together with the inverse action of

the Jacobian (useful in design and optimization) at convergence. We have demonstrated NKS methods on realistic 3D NASA flow configurations (using the NASA code FUN3D in both incompressible and compressible regimes) and gained experience tuning their many parameters so that analysts and optimizers can be attracted to their use. High efficiencies can be obtained; however, in scaling NKS methods to the largest systems conveniently available to us presently (512 processors) and solving unstructured problems with over 10 million unknowns we have observed, without surprise, that global communication from norm and inner product computations is the major detractor from linear speedup. The majority of these global operations arise in solving the Newton correction equations (large sparse linear systems) by Krylov methods. Our recent efforts have therefore concentrated on algorithmic adaptations of NKS appropriate for fine-grained parallelism, such as will be required at teraflop/s scales, to reduce the time consumed by these global operations.

There are two means of reducing the time spent on nonlocal operations: making them less frequent and making them less costly when they must occur.

With respect to the latter, we have observed that the major cost of nonlocal operations is in their synchronization, not in the actual transfer or manipulation of data. Synchronization exposes load imbalance. Our “volume” work, the dominant term in processor work estimates, is nearly perfectly balanced (we use publicly available partitioners, like McTiS and Jostle). However, our “surface” (interface) work is imbalanced. This is not an issue with relatively coarse granularity (e.g., less than 100 processors for a million-degree of freedom simulation) but “surface” work is increasingly important when the number of degrees of freedom per processor slips below 10,000 (a number that depends on problem and architecture, but would be typical for implicit unstructured CFD on platforms in the IBM SP and Cray T3 series).

With respect to the former, we reserve the necessity of full NKS as final convergence is approached, in order to derive the asymptotic convergence rate benefits of Newton’s method, but we can replace early global Newton corrections with separate regional Newton corrections in a family of algorithms we refer to as “nonlinear Schwarz.” The regional applications of NKS need to synchronize only internally, with other processors sharing the region. They can either be completely asynchronous with respect to other regions, or be loosely synchronous. This should extend useful parallel granularity in implicit parallel CFD to arbitrarily fine levels, provided that the overall nonlinear convergence rate does not suffer. We have been exploring nonlinear Schwarz on a model nonlinear diffusion problem possessing a boundary layer induced by the constitutive law. (This problem is prototypical of CFD and many other real-world applications in the sense that the strong nonlinearity is confined to a relatively small region, and the majority of the domain has linear dynamics. For instance, the same phenomenon was observed in transonic full potential flow over the NACA0012, where the formation of the shock controls the global Newton convergence though the updates are significant only in the near-wing region.) We find that the global problem can be solved in just a few global NKS iterations once the regional problem containing the strong nonlinearity is internally converged. In addition to demonstrating the scalability of nonlinear Schwarz in this context, this finding opens up a vast new frontier for algorithmic adaptivity based on local degree of nonlinearity.

We will continue to pioneer nonlinear Schwarz methods in implicit parallel CFD, including developing metrics for nonlinearity and developing a software environment appropriate for less synchronous computation. We will work on improved partitioning methods, in order to balance the sum of volume and surface work, rather than just the volume work, as at present. We will also work on partitioning methods that preserve strong flow-oriented couplings, for improved convergence rates. We will also continue to work with the CFD user community to make workstation network and HPCCP testbed use “friendly.” This includes attention to parallel I/O, checkpointing, and mesh sequencing. Finally, we will explore the additional task parallelism

available in use of the parallel analysis codes in optimization.

This work was performed in collaboration with W.K. Anderson (NASA Langley), D. Kaushik and N. Karunaratne (Old Dominion University), X.-C. Cai (UC-Boulder), and W.D. Gropp, L.C. McInnes, and B.F. Smith (Argonne National Laboratory).

KWAN-LIU MA

Efficient Encoding and Rendering of Time-Varying Volume Data

Visualization of time-varying volumetric data sets, which may be obtained from numerical simulations or sensing instruments, provides scientists insights into the detailed dynamics of the phenomenon under study. This research seeks a coherent solution based on quantization, coupled with octree and difference encoding for visualizing time-varying volumetric data. Our goal is to increase the users interaction with the data. This requires that the images be presented to the user as rapidly as possible. Although we do not see large savings when the cost of encoding and rendering are combined, by preprocessing we can achieve near interactive viewing rates.

Quantization is used to attain voxel-level compression and may have a significant influence on the performance of the subsequent encoding and visualization steps. Octree encoding is used for spatial domain compression, and difference encoding for temporal domain compression. In essence, neighboring voxels may be fused into macro voxels if they have similar values, and subtrees at consecutive time steps may be merged if they are identical. The software rendering process is tailored according to the tree structures and the volume visualization process. With the tree representation, selective rendering may be performed very efficiently. Additionally, the I/O costs are reduced. With these combined savings, a higher level of user interactivity is achieved. We have studied a variety of time-varying volume datasets, performed encoding based on data statistics, and optimized the rendering calculations wherever possible. We found that, in general, the selection of encoding and rendering strategies should depend very much on data resolution, statistics and visualization requirements. The savings in storage space also reduces the I/O required by the renderer. With large datasets over long intervals of time, this reduction can be a significant part of the overall savings. Preliminary tests on workstations have shown in many cases tremendous reduction by as high as 90% in both storage space and inter-frame delay.

Future work includes the development of application-specific techniques and taking the grid structures (curvilinear, unstructured, etc.) into consideration. We will investigate how the order of encoding calculations would impact the overall compression and rendering performance. In addition, we will study the characteristics of time-varying computational fluid dynamics datasets and continue developing appropriate compression and rendering methods.

PIYUSH MEHROTRA

Arcade: A Distributed Computing Environment for ICASE

Distributed heterogeneous computing is being increasingly applied to a variety of large size computational problems. Such computations, for example, the multidisciplinary design optimization of an aircraft, generally consist of multiple heterogeneous modules interacting with each other to solve the problem at hand. Such applications are generally developed by a team in which each discipline is the responsibility of experts in the field. The objective of this project is to develop a GUI-based environment which supports the multi-user

design of such applications and their execution and monitoring in a heterogeneous environment consisting of a network of workstations, specialized machines and parallel architectures.

The overall goal is to design an interface which is easy to use, easily accessible, and portable. We are planning to leverage off of technologies where available to achieve these goals. In particular, integrating the interface in a web browser, e.g., Netscape, provides users with a familiar interface on desktops ranging from Unix based workstations, to Windows based PCs, and Macintoshes. The Arcade system consists of the following sub-environments: *application design interface*, *resource allocation and execution interface* and *monitoring and steering interface*. The architecture is a three-tier architecture. The front-end, based on Java, provides the user interface. The middle tier consists of the user interface server which provides the logic to process the user input and to interact with application modules running on a heterogeneous set of machines. Application execution is monitored by the Execution Controller which interacts Process Controllers on the individual resources. The PCs along with application modules form the last tier of the architecture.

The current implementation is focusing on the visual interface for specifying the modules and their interactions. The interface allows the user to specify hierarchical modules, control structures such as conditionals and loops, and SPMD modules to be run on multiple processors using MPI for communication. The graphical representation of the application is also used to indicate the current execution status of the application, i.e., which modules have finished execution, which modules are currently executing, and which are waiting for execution.

We are currently expanding the system to incorporate all the facilities envisioned in our architecture in a phased approach, an approach where the driving force is the user of the system. In particular, we are focusing on the support required for multi-domain execution of the modules. The critical issue here is the security and authentication required for an application launched in one administrative domain to execute part of its code in another domain.

This work was performed in collaboration with Z. Chen, K. Maly, A. Al-Theneyan, M. Zubair (Old Dominion University).

Multithreaded System for Distributed Environments

Traditionally, lightweight threads are supported only within the single address space of a process, or in shared memory environments with multiple processes. Likewise, interprocess communication systems do not currently allow messages to be sent directly to entities within a process. The objective of this project is to build a system which combines standard interfaces for lightweight threads, pthreads, and interprocess communication, MPI, to support point-to-point communication between any two threads in a distributed memory system.

The Chant runtime system has been built using layers: point-to-point communication, remote service requests, remote thread operations. In the last year we have added a layer to support load balancing via migration of threads. In contrast to other thread migration systems, we provide migration in the presence of pointers, i.e., along with the stack, the thread heap is also migrated. This allows the pointers to point to valid data even after the migration. The load balancing layer provides facilities for computing the current workloads of the processors, figuring which threads should be moved and to what processors. The system also keeps track of the amount of communication being generated by each thread so that it can be used in the decision making process. The underlying thread migration layer takes care of the actual motion of the threads. The system is designed such that default routines can be replaced by the user with specialized routines for carrying out the various functions.

We have tested the system on several applications including a simple adaptive quadrature code, a traveling salesman code, and an image rendering code. In each of these cases we have shown that there is an overall gain in performance due to thread-based load balancing. That is, the overhead due to multi-threading, thread migration and load management, as implemented in our system are at acceptable levels.

We are continuing to test and optimize the load-balancing layer via several thread-based parallel codes.

This work was performed in collaboration with D. Cronk (The College of William & Mary).

Evaluation and Extension of High Performance Fortran

The stated goal of High Performance Fortran (HPF) was to “address the problems of writing data parallel programs where the distribution of data affects performance.” We have been using data parallel codes of interest to NASA to evaluate the effectiveness of the features of HPF and to suggest extensions for the language.

We have continued our exploration of using HPF for unstructured grid codes. These codes use indexing arrays to access the data and thus the data access pattern is not known at compile-time. HPF compilers generally use an inspector-executor strategy to analyze and generate the communication schedule at runtime in the inspector phase and use it in the executor phase to affect the actual communication. The inspector phase is generally expensive and needs to be amortized over several repetitions of the executor. We have devised directives that allow users to specify when to reuse the schedule and when it is necessary to recompute the schedule. These directives were implemented in the Vienna Fortran Compiler being developed at the University of Vienna. Tests run on several codes. Both benchmarks and kernels from real applications show that using these directives leads to substantial performance gains.

Recently a proposal was put forth for a set of language extensions to Fortran and C based upon a fork-join model of parallel execution; called OpenMP, it aims to provide a portable shared memory programming interface for shared memory and low latency systems. However, these extensions ignore the issue of data locality which becomes a performance issue on shared address space machines which use a physically distributed memory system. We have been investigating how the two models, HPF and OpenMP, can be used together to write programs which exploit the full capabilities of such systems. One approach is to use the HPF extrinsic procedures to call OpenMP routines. This allows both models to be separately used. A more integrated approach is to provide HPF-like mapping directives in OpenMP so as to be able to control the data locality. We are currently investigating both of these approaches for several different applications. This work is being done in collaboration with B. Chapman (Univ. of Vienna).

A new version of HPF (HPF 2.0) which includes several extensions resulting from our experiments, has been released. We plan to continue the evaluation of the new features of HPF as they are incorporated in the compilers.

This work is being done in collaboration with K. Roe (The College of William & Mary), J. Van Rosendale (ICASE & NSF), and H. Zima (University of Vienna).

ALEX POTHEN

Object-Oriented Design of Sparse Solvers

Object-oriented design (OOD) is a paradigm for managing the increasing complexity of software; it also helps to make software flexible, extensible, and easier to maintain. We have employed object-oriented design to create software for solving sparse systems of linear equations by direct methods. Direct methods employ

sophisticated combinatorial and algebraic algorithms that contribute to software complexity, and hence it is natural to consider OOD in this context.

OOD manages complexity by means of decomposition and abstraction. We decompose our software into two main types of objects: structural objects corresponding to data structures and algorithmic objects corresponding to algorithms. This design decouples data structures from algorithms, permitting a user to experiment with different algorithms and different data structures, and, if necessary, develop new algorithms and data structures. However, this increased flexibility comes at the cost of some loss in efficiency. We have made careful trade-offs in our software to achieve the benefits of OOD without sacrificing efficiency. The running times of our C++ code for the symmetric positive definite solver compare quite favorably with existing Fortran 77 codes. We have used our code to solve Helmholtz problems (complex symmetric) from electromagnetics, and symmetric indefinite problems from structural analysis, fluid dynamics, and linear programming. This is the first object-oriented implementation of sparse direct methods that we know of.

We are continuing to test and develop our solvers. The next step is the creation of parallel and out-of-core solvers. We will share the current version of the code with users.

This work was performed in collaboration with F. Dobrian and G. Kumfert (Old Dominion University).

Fast Algorithms for Incomplete Factorization Preconditioners

The parallel computation of robust preconditioners is a priority for solving large systems of equations on unstructured grids and in other applications. We have developed new algorithms and software that can compute incomplete factorization preconditioners for high level fill in time proportional to the number of floating point operations and memory accesses.

We have developed a structure theory based on paths in the adjacency graph of the matrix to predict where zero elements become nonzeros in incomplete factorization (fill elements). A level function is used in incomplete factorization to control the number of fill elements, and we relate the level of fill to lengths of appropriately defined paths in the adjacency graph. This result permits us to search in the neighborhood of a vertex in the graph to predict all fill elements associated with that vertex. We have designed two variants of these algorithms and have proven that they have a smaller running time complexity than currently used algorithms for computing incomplete factorizations. The more efficient algorithms make use of the concept of transitive reduction of directed graphs (symmetric problems) and symmetric reduction of directed graphs (unsymmetric problems) in order to search for paths in smaller graphs. Our current implementation in C shows that the new algorithms are faster than the currently used implementations. We presented our results at the Copper Mountain Conference on Iterative Methods in April.

We will continue to develop algorithms and software for fast computation of preconditioners, and investigate parallel implementations. The parallel algorithms will make it possible to compute high level incomplete factorization preconditioners for parallel Krylov-space-based iterative solvers.

This work was performed in collaboration with D. Hysom (Old Dominion University, who is also affiliated with ICASE).

REPORTS AND ABSTRACTS

Alexandrov, Natalia, J.E. Dennis, Jr., Robert Michael Lewis, and Virginia Torczon: *A trust region framework for managing the use of approximation models in optimization*. ICASE Report No. 97-50, (NASA CR-201745), October 2, 1997, 19 pages. Submitted to Structural Optimization.

This paper presents an analytically robust, globally convergent approach to managing the use of approximation models of various fidelity in optimization. By robust global behavior we mean the mathematical assurance that the iterates produced by the optimization algorithm, started at an arbitrary initial iterate, will converge to a stationary point or local optimizer for the original problem. The approach we present is based on the trust region idea from nonlinear programming and is shown to be provably convergent to a solution of the original high-fidelity problem. The proposed method for managing approximations in engineering optimization suggests ways to decide when the fidelity, and thus the cost, of the approximations might be fruitfully increased or decreased in the course of the optimization iterations. The approach is quite general. We make no assumptions on the structure of the original problem, in particular, no assumptions of convexity and separability, and place only mild requirements on the approximations. The approximations used in the framework can be of any nature appropriate to an application; for instance, they can be represented by analyses, simulations, or simple algebraic models. This paper introduces the approach and outlines the convergence analysis.

Guattery, Stephen, Tom Leighton, and Gary L. Miller: *The path resistance method for bounding the smallest nontrivial eigenvalue of a Laplacian*. ICASE Report No. 97-51, (NASA CR-201746), October 6, 1997, 20 pages. Submitted to Combinatorics, Probability, and Computing.

We introduce the path resistance method for lower bounds on the smallest nontrivial eigenvalue of the Laplacian matrix of a graph. The method is based on viewing the graph in terms of electrical circuits; it uses clique embeddings to produce lower bounds on λ_2 and star embeddings to produce lower bounds on the smallest Rayleigh quotient when there is a zero Dirichlet boundary condition. The method assigns priorities to the paths in the embedding; we show that, for an unweighted tree T , using uniform priorities for a clique embedding produces a lower bound on λ_2 that is off by at most an $O(\log \text{diameter}(T))$ factor. We show that the best bounds this method can produce for clique embeddings are the same as for a related method that uses clique embeddings and edge lengths to produce bounds.

Hesthaven, J.S., and L.M. Jameson: *A wavelet optimized adaptive multi-domain method*. ICASE Report No. 97-52, (NASA CR-201747), October 8, 1997, 21 pages. Submitted to Journal of Computational Physics.

The formulation and implementation of wavelet based methods for the solution of multi-dimensional partial differential equations in complex geometries is discussed. Utilizing the close connection between Daubechies wavelets and finite difference methods on arbitrary grids, we formulate a wavelet based collocation method, well suited for dealing with general boundary conditions and nonlinearities.

To circumvent problems associated with completely arbitrary grids and complex geometries, we propose to use a multi-domain formulation to solve the partial differential equation, with the ability to adapt the

grid as well as the order of the scheme within each subdomain. In addition to supplying the required geometric flexibility, the multi-domain formulation also provides a very natural load-balanced data-decomposition, suitable for parallel environments.

The performance of the overall scheme is illustrated by solving two-dimensional hyperbolic problems.

Ristorcelli, J.R., and G.A. Blaisdell: *Validation of a pseudo-sound theory for the pressure-dilatation in DNS of compressible turbulence*. ICASE Report No. 97-53, (NASA CR-201748), October 16, 1997, 14 pages. To appear in the Proceedings of the 11th Turbulent Shear Flow Symposium, Grenoble, France.

The results of an asymptotic theory for statistical closures for compressible turbulence are explored and validated with the direct numerical simulation of the isotropic decay and the homogeneous shear. An excellent collapse of the data is seen. The slow portion is found to scale, as predicted by the theory, with the quantity M_t^2 and ε_s . The rapid portion has an unambiguous scaling with $\alpha^2 M_t^2 \varepsilon_s [P_k/\varepsilon - 1] (Sk/\varepsilon)^2$. Implicit in the scaling is a dependence, as has been noted by others, on the gradient Mach number. A new feature of the effects of compressibility, that of the Kolmogorov scaling coefficient, α , is discussed. It is suggested that α may contain flow specific physics associated with large scales that might provide further insight into the structural effects of compressibility.

Girimaji, Sharath S.: *Dynamical system analysis of Reynolds stress closure equations*. ICASE Report No. 97-54, (NASA CR-201749), October 17, 1997, 17 pages. To appear in the Proceedings of the 11th Turbulent Shear Flow Symposium, Grenoble, France.

In this paper, we establish the causality between the model coefficients in the standard pressure-strain correlation model and the predicted equilibrium states for homogeneous turbulence. We accomplish this by performing a comprehensive fixed point analysis of the modeled Reynolds stress and dissipation rate equations. The results from this analysis will be very useful for developing improved pressure-strain correlation models to yield observed equilibrium behavior.

Allan, Brian G., Maurice Holt, and Andrew Packard: *Simulation of a controlled airfoil with jets*. ICASE Report No. 97-55, (NASA CR-201750), October 16, 1997, 21 pages. To appear in the Journal of Guidance, Control, and Dynamics.

Numerical simulations of a two-dimensional airfoil, controlled by an applied moment in pitch and an airfoil controlled by jets, were investigated. These simulations couple the Reynolds-averaged Navier-Stokes equations and Euler's equations of rigid body motion, with an active control system. Controllers for both systems were designed to track altitude commands and were evaluated by simulating a closed-loop altitude step response using the coupled system. The airfoil controlled by a pitching moment used an optimal state feedback controller. A closed-loop simulation, of the airfoil with an applied moment, showed that the trajectories compared very well with quasi-steady aerodynamic theory, providing a measure of validation. The airfoil with jets used a controller designed by robust control methods. A linear plant model for this system was identified using open-loop data generated by the nonlinear coupled system. A closed-loop simulation of the airfoil with jets, showed good tracking of an altitude command. This simulation also showed oscillations in the control input as a result of dynamics not accounted for in the control design. This research work demonstrates how computational fluid dynamics, coupled with rigid body dynamics, and a control law can be used to prototype control systems in problematic nonlinear flight regimes.

Banerjee, Nana S., and James F. Geer: *Exponential approximations using Fourier series partial sums*. ICASE Report No. 97-56, (NASA CR-201751), October 20, 1997, 42 pages. Submitted to SIAM Journal on Scientific Computation.

The problem of accurately reconstructing a piece-wise smooth, 2π -periodic function f and its first few derivatives, given only a truncated Fourier series representation of f , is studied and solved. The reconstruction process is divided into two steps. In the first step, the first $2N + 1$ Fourier coefficients of f are used to approximate the locations and magnitudes of the discontinuities in f and its first M derivatives. This is accomplished by first finding initial estimates of these quantities based on certain properties of Gibbs phenomenon, and then refining these estimates by fitting the asymptotic form of the Fourier coefficients to the given coefficients using a least-squares approach. It is conjectured that the locations of the singularities are approximated to within $O(N^{-M-2})$, and the associated jump of the k^{th} derivative of f is approximated to within $O(N^{-M-1+k})$, as $N \rightarrow \infty$, and the method is robust. These estimates are then used with a class of singular basis functions, which have certain “built-in” singularities, to construct a new sequence of approximations to f . Each of these new approximations is the sum of a piecewise smooth function and a new Fourier series partial sum. When N is proportional to M , it is shown that these new approximations, and their derivatives, converge exponentially in the maximum norm to f , and its corresponding derivatives, except in the union of a finite number of small open intervals containing the points of singularity of f . The total measure of these intervals decreases exponentially to zero as $M \rightarrow \infty$. The technique is illustrated with several examples.

Arian, Eyal: *Convergence estimates for multidisciplinary analysis and optimization*. ICASE Report No. 97-57, (NASA CR-201752), October 17, 1997, 39 pages. To appear in Progress in Systems and Control Theory (Birkhäuser Volume).

A quantitative analysis of coupling between systems of equations is introduced. This analysis is then applied to problems in multidisciplinary analysis, sensitivity, and optimization. For the sensitivity and optimization problems both multidisciplinary and single discipline feasibility schemes are considered. In all these cases a “convergence factor” is estimated in terms of the Jacobians and Hessians of the system, thus it can also be approximated by existing disciplinary analysis and optimization codes.

The convergence factor is identified with the measure for the “coupling” between the disciplines in the system.

Applications to algorithm development are discussed. Demonstration of the convergence estimates and numerical results are given for a system composed of two non-linear algebraic equations, and for a system composed of two PDEs modeling aeroelasticity.

Girimaji, S.S., and J.R. Ristorcelli: *On the behavior of velocity fluctuations in rapidly rotating flows*. ICASE Report No. 97-58, (NASA CR-97-206244), November 17, 1997, 21 pages. To be submitted to Journal of Fluid Mechanics.

The behavior of velocity fluctuations subjected to rapid rotation is examined. The rapid rotation considered is any arbitrary combination of two basic forms of rotation, reference frame rotation and mean flow rotation. It is recognized that the two types of rotating flows differ in the manner in which the fluctuating fields are advected. The first category is comprised of flows in rotating systems of which synoptic scale geophysical flows are a good example. In this class of flows the fluctuating velocity field advects and rotates

with the mean flow. In the rapid rotation limit, the Taylor-Proudman theorem describes the behavior of this class of fluctuations. Velocity fluctuations that are advected without rotation by the mean flow constitute the second category which includes vortical flows of aerodynamic interest. The Taylor-Proudman theorem is not pertinent to this class flows and a new result appropriate to this second category of fluctuations is derived. The present development demonstrates that the fluctuating velocity fields are rendered two-dimensional and horizontally non-divergent in the limit of any large combination of reference frame rotation and mean-flow rotation. The concomitant ‘geostrophic’ balance of the momentum equation is, however, dependent upon the form of rapid rotation. It is also demonstrated that the evolution equations of a two-dimensional fluctuating velocity fields are frame-indifferent with any imposed mean-flow rotation. The analyses and results of this paper highlight many fundamental aspects of rotating flows and have important consequences for their turbulence closures in inertial and non-inertial frames.

del Rosario, R.C.H., and R.C. Smith: *LQR control of thin shell dynamics: Formulation and numerical implementation*. ICASE Report No. 97-59, (NASA CR-97-206245), November 13, 1997, 33 pages. Submitted to Journal of Intelligent Material Systems and Structures.

A PDE-based feedback control method for thin cylindrical shells with surface-mounted piezoceramic actuators is presented. Donnell-Mushtari equations modified to incorporate both passive and active piezoceramic patch contributions are used to model the system dynamics. The well-posedness of this model and the associated LQR problem with an unbounded input operator are established through analytic semigroup theory. The model is discretized using a Galerkin expansion with basis functions constructed from Fourier polynomials tensored with cubic splines, and convergence criteria for the associated approximate LQR problem are established. The effectiveness of the method for attenuating the coupled longitudinal, circumferential and transverse shell displacements is illustrated through a set of numerical examples.

Calkins, F.T., R.C. Smith, and A.B. Flatau: *An energy-based hysteresis model for magnetostrictive transducers*. ICASE Report No. 97-60, (NASA CR-97-206246), November 17, 1997, 21 pages. Submitted to IEEE Transactions on Magnetics.

This paper addresses the modeling of hysteresis in magnetostrictive transducers. This is considered in the context of control applications which require an accurate characterization of the relation between input currents and strains output by the transducer. This relation typically exhibits significant nonlinearities and hysteresis due to inherent properties of magnetostrictive materials. The characterization considered here is based upon the Jiles-Atherton mean field model for ferromagnetic hysteresis in combination with a quadratic moment rotation model for magnetostriction. As demonstrated through comparison with experimental data, the magnetization model very adequately quantifies both major and minor loops under various operating conditions. The combined model can then be used to accurately characterize output strains at moderate drive levels. The advantages to this model lie in the small number (six) of required parameters and the flexibility it exhibits in a variety of operating conditions.

Lewis, Robert Michael: *Numerical computation of sensitivities and the adjoint approach*. ICASE Report No. 97-61, (NASA CR-97-206247), November 17, 1997, 19 pages. Submitted to the Proceedings of the AFOSR Workshop on Optimal Design, Arlington, VA (Sept. 30 - Oct. 3, 1997).

We discuss the numerical computation of sensitivities via the adjoint approach in optimization problems governed by differential equations. We focus on the adjoint problem in its weak form. We show how one can

avoid some of the problems with the adjoint approach, such as deriving suitable boundary conditions for the adjoint equation. We discuss the convergence of numerical approximations of the costate computed via the weak form of the adjoint problem and show the significance for the discrete adjoint problem.

Bataille, F., Ye Zhou, and Jean-Pierre Bertoglio: *Energy transfer and triadic interactions in compressible turbulence*. ICASE Report No. 97-62, (NASA CR-97-206249), November 19, 1997, 30 pages. To be submitted to Journal of Fluid Mechanics.

Using a two-point closure theory, the Eddy-Damped-Quasi-Normal-Markovian (EDQNM) approximation, we have investigated the energy transfer process and triadic interactions of compressible turbulence. In order to analyze the compressible mode directly, the Helmholtz decomposition is used. The following issues were addressed: (1) What is the mechanism of energy exchange between the solenoidal and compressible modes, and (2) Is there an energy cascade in the compressible energy transfer process? It is concluded that the compressible energy is transferred locally from the solenoidal part to the compressible part. It is also found that there is an energy cascade of the compressible mode for high turbulent Mach number ($M_t \geq 0.5$). Since we assume that the compressibility is weak, the magnitude of the compressible (radiative or cascade) transfer is much smaller than that of solenoidal cascade. These results are further confirmed by studying the triadic energy transfer function, the most fundamental building block of the energy transfer.

Rubinstein, Robert, and Ye Zhou: *The dissipation rate transport equation and subgrid-scale models in rotating turbulence*. ICASE Report No. 97-63, (NASA CR-97-206250), November 19, 1997, 17 pages. To appear in Proceedings of the 11th Turbulent Shear Flow Conference.

The dissipation rate transport equation remains the most uncertain part of turbulence modeling. The difficulties are increased when external agencies like rotation prevent straightforward dimensional analysis from determining the correct form of the modelled equation. In this work, the dissipation rate transport equation and subgrid scale models for rotating turbulence are derived from an analytical statistical theory of rotating turbulence. In the strong rotation limit, the theory predicts a turbulent steady state in which the inertial range energy spectrum scales as k^{-2} and the turbulent time scale is the inverse rotation rate. This scaling has been derived previously by heuristic arguments.

Yeung, P.K., and Ye Zhou: *On the universality of the Kolmogorov constant in numerical simulations of turbulence*. ICASE Report No. 97-64, (NASA CR-97-206251), November 21, 1997, 18 pages. Submitted to Physical Review E.

Motivated by a recent survey of experimental data [K.R. Sreenivasan, Phys. Fluids **7**, 2778 (1995)], we examine data on the Kolmogorov spectrum constant in numerical simulations of isotropic turbulence, using results both from previous studies and from new direct numerical simulations over a range of Reynolds numbers (up to 240 on the Taylor scale) at grid resolutions up to 512^3 . It is noted that in addition to $k^{-5/3}$ scaling, identification of a true inertial range requires spectral isotropy in the same wavenumber range. We found that a plateau in the compensated three-dimensional energy spectrum at $k\eta \approx 0.1 - 0.2$, commonly used to infer the Kolmogorov constant from the compensated three-dimensional energy spectrum, actually does not represent proper inertial range behavior. Rather, a proper, if still approximate, inertial range emerges at $k\eta \approx 0.02 - 0.05$ when R_λ increases beyond 140. The new simulations indicate proportionality constants C_1 and C in the one- and three-dimensional energy spectra respectively about 0.60 and 1.62. If

the turbulence were perfectly isotropic then use of isotropy relations in wavenumber space ($C_1 = 18/55 C$) would imply that $C_1 \approx 0.53$ for $C = 1.62$, in excellent agreement with experiments. However the one- and three-dimensional estimates are not fully consistent, because of departures (due to numerical and statistical limitations) from isotropy of the computed spectra at low wavenumbers. The inertial scaling of structure functions in physical space is briefly addressed. Since DNS is still restricted to moderate Reynolds numbers, an accurate evaluation of the Kolmogorov constant is very difficult. We focus on providing new insights on the interpretation of Kolmogorov 1941 similarity in the DNS literature and do not consider issues pertaining to the refined similarity hypotheses of Kolmogorov (K62).

Shu, Chi-Wang: *Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws*. ICASE Report No. 97-65, (NASA CR-97-206253), November 21, 1997, 83 pages. To appear in Lecture Notes in Mathematics, CIME Subseries, Springer Verlag.

In these lecture notes we describe the construction, analysis, and application of ENO (Essentially Non-Oscillatory) and WENO (Weighted Essentially Non-Oscillatory) schemes for hyperbolic conservation laws and related Hamilton-Jacobi equations. ENO and WENO schemes are high order accurate finite difference schemes designed for problems with piecewise smooth solutions containing discontinuities. The key idea lies at the approximation level, where a nonlinear adaptive procedure is used to automatically choose the locally smoothest stencil, hence avoiding crossing discontinuities in the interpolation procedure as much as possible. ENO and WENO schemes have been quite successful in applications, especially for problems containing both shocks and complicated smooth solution structures, such as compressible turbulence simulations and aeroacoustics.

These lecture notes are basically self-contained. It is our hope that with these notes and with the help of the quoted references, the reader can understand the algorithms and code them up for applications. Sample codes are also available from the author.

Buchholz, Peter, Gianfranco Ciardo, Susanna Donatelli, and Peter Kemper: *Complexity of Kronecker operations on sparse matrices with applications to the solution of Markov models*. ICASE Report No. 97-66, (NASA CR-97-206274), December 18, 1997, 29 pages. To be submitted to INFORMS Journal on Computing.

We present a systematic discussion of algorithms to multiply a vector by a matrix expressed as the Kronecker product of sparse matrices, extending previous work in a unified notational framework. Then, we use our results to define new algorithms for the solution of large structured Markov models. In addition to a comprehensive overview of existing approaches, we give new results with respect to: (1) managing certain types of state-dependent behavior without incurring extra cost; (2) supporting both Jacobi-style and Gauss-Seidel-style methods by appropriate multiplication algorithms; (3) speeding up algorithms that consider probability vectors of size equal to the “actual” state space instead of the “potential” state space.

Ciardo, Gianfranco, and Marco Tilgner: *Parametric state space structuring*. ICASE Report No. 97-67, (NASA CR-97-206267), December 10, 1997, 30 pages. Submitted to IEEE Transactions on Software Engineering.

Structured approaches based on Kronecker operators for the description and solution of the infinitesimal generator of a continuous-time Markov chains are receiving increasing interest. However, their main advantage, a substantial reduction in the memory requirements during the numerical solution, comes at a

price. Methods based on the “potential state space” allocate a probability vector that might be much larger than actually needed. Methods based on the “actual state space”, instead, have an additional logarithmic overhead. We present an approach that realizes the advantages of both methods with none of their disadvantages, by partitioning the local state spaces of each submodel. We apply our results to a model of software rendezvous, and show how they reduce memory requirements while, at the same time, improving the efficiency of the computation.

Jameson, Antony: *Essential elements of computational algorithms for aerodynamic analysis and design*. ICASE Report No. 97-68, (NASA CR-97-206268), December 10, 1997, 71 pages.

This paper traces the development of computational fluid dynamics as a tool for aircraft design. It addresses the requirements for effective industrial use, and trade-offs between modeling accuracy and computational costs. Essential elements of algorithm design are discussed in detail, together with a unified approach to the design of shock capturing schemes. Finally, the paper discusses the use of techniques drawn from control theory to determine optimal aerodynamic shapes. In the future multidisciplinary analysis and optimization should be combined to provide an integrated design environment.

Arian, Eyal, and Manuel D. Salas: *Admitting the inadmissible: Adjoint formulation for incomplete cost functionals in aerodynamic optimization*. ICASE Report No. 97-69, (NASA CR-97-206269), December 10, 1997, 18 pages. To be submitted to AIAA Journal.

We derive the adjoint equations for problems in aerodynamic optimization which are improperly considered as “inadmissible”. For example, a cost functional which depends on the density, rather than on the pressure, is considered “inadmissible” for an optimization problem governed by the Euler equations. We show that for such problems additional terms should be included in the Lagrangian functional when deriving the adjoint equations. These terms are obtained from the restriction of the interior PDE to the control surface. Demonstrations of the explicit derivation of the adjoint equations for “inadmissible” cost functionals are given for the potential, Euler, and Navier-Stokes equations.

Chieuh, Tzi-Cker, and Kwan-Liu Ma: *A parallel pipelined renderer for time-varying volume data*. ICASE Report No. 97-70, (NASA CR-97-206275), December 18, 1997, 17 pages. Submitted to the International Journal in High Performance Computer Graphics, Multimedia and Visualisation.

This paper presents a strategy for efficiently rendering time-varying volume data sets on a distributed-memory parallel computer. Time-varying volume data take large storage space and visualizing them requires reading large files continuously or periodically throughout the course of the visualization process. Instead of using all the processors to collectively render one volume at a time, a pipelined rendering process is formed by partitioning processors into groups to render multiple volumes concurrently. In this way, the overall rendering time may be greatly reduced because the pipelined rendering tasks are overlapped with the I/O required to load each volume into a group of processors; moreover, parallelization overhead may be reduced as a result of partitioning the processors. We modify an existing parallel volume renderer to exploit various levels of rendering parallelism and to study how the partitioning of processors may lead to optimal rendering performance. Two factors which are important to the overall execution time are resource utilization efficiency and pipeline startup latency. The optimal partitioning configuration is the one that

balances these two factors. Tests on Intel Paragon computers show that in general optimal partitionings do exist for a given rendering task and result in 40-50% saving in overall rendering time.

Soemarwoto, Bambang: *The variational method for aerodynamic optimization using the Navier-Stokes equations*. ICASE Report No. 97-71, (NASA CR-97-206277), December 19, 1997, 32 pages. To be submitted to Computer and Fluids.

This report describes the formulation of an aerodynamic shape design methodology using a compressible viscous flow model based on the Reynolds-Averaged Navier-Stokes equations. The aerodynamic shape is described by a set of geometrical design variables. The design problem is formulated as an optimization problem stated in terms of an aerodynamic objective functional which has to be minimized. The design scheme employs a gradient-based optimization algorithm in order to obtain the optimum values of the design variables. The gradient of the aerodynamic functional with respect to the design variables is computed by means of the variational method, which requires the solution of an adjoint problem. The formulation of the adjoint problem is described which leads to a set of adjoint equations and boundary conditions. Using the flow variables and the adjoint variables, an expression for the gradient has been constructed. Computational results are presented for an inverse problem of an airfoil. It will be shown that, starting from an initial geometry of the NACA 0012 airfoil, the target pressure distribution and geometry of a best-fit of the RAE 2822 airfoil in a transonic flow condition has been reconstructed successfully.

Hayder, M. Ehtesham, David E. Keyes, and Piyush Mehrotra: *A comparison of PETSc library and HPF implementation of an archetypal PDE computation*. ICASE Report No. 97-72, (NASA CR-97-206278), December 19, 1997, 27 pages. Submitted to Journal for Advances in Engineering Software.

Two paradigms for distributed-memory parallel computation that free the application programmer from the details of message passing are compared for an archetypal structured scientific computation — a non-linear, structured-grid partial differential equation boundary value problem — using the same algorithm on the same hardware. Both paradigms, parallel libraries represented by Argonne's PETSc, and parallel languages represented by the Portland Group's HPF, are found to be easy to use for this problem class, and both are reasonably effective in exploiting concurrency after a short learning curve. The level of involvement required by the application programmer under either paradigm includes specification of the data partitioning (corresponding to a geometrically simple decomposition of the domain of the PDE). Programming in SPMD style for the PETSc library requires writing the routines that discretize the PDE and its Jacobian, managing subdomain-to-processor mappings (affine global-to-local index mappings), and interfacing to library solver routines. Programming for HPF requires a complete sequential implementation of the same algorithm, introducing concurrency through subdomain blocking (an effort similar to the index mapping), and modest experimentation with rewriting loops to elucidate to the compiler the latent concurrency. Correctness and scalability are cross-validated on up to 32 nodes of an IBM SP2.

Keyes, David E., Dinesh K. Kaushik, and Barry F. Smith: *Prospects for CFD on petaflops systems*. ICASE Report No. 97-73, (NASA CR-97-206279), December 19, 1997, 30 pages. Submitted for CFD Review 1997.

With teraflops-scale computational modeling expected to be routine by 2003-04, under the terms of the Accelerated Strategic Computing Initiative (ASCI) of the U.S. Department of Energy, and with teraflops-

capable platforms already available to a small group of users, attention naturally focuses on the next symbolically important milestone, computing at rates of 10^{15} floating point operations per second, or “petaflop/s”. For architectural designs that are in any sense extrapolations of today’s, petaflops-scale computing will require approximately one-million-fold instruction-level concurrency. Given that cost-effective one-thousand-fold concurrency is challenging in practical computational fluid dynamics simulations today, algorithms are among the many possible bottlenecks to CFD on petaflops systems. After a general outline of the problems and prospects of petaflops computing, we examine the issue of algorithms for PDE computations in particular. A back-of-the-envelope parallel complexity analysis focuses on the latency of global synchronization steps in the implicit algorithm. We argue that the latency of synchronization steps is a fundamental, but addressable, challenge for PDE computations with static data structures, which are primarily determined by grids. We provide recent results with encouraging scalability for parallel implicit Euler simulations using the Newton-Krylov-Schwarz solver in the PETSc software library. The prospects for PDE simulations with dynamically evolving data structures are far less clear.

Ristorcelli, J.R.: *Some results relevant to statistical closures for compressible turbulence*. ICASE Report No. 98-1, (NASA CR-1998-206902), January 23, 1998, 28 pages. To appear in the Proceedings from “Modeling Complex Turbulent Flows,” by Kluwer Academic Publishers.

For weakly compressible turbulent fluctuations there exists a small parameter, the square of the fluctuating Mach number, that allows an investigation using a perturbative treatment. The consequences of such a perturbative analysis in three different subject areas are described: 1) initial conditions in direct numerical simulations, 2) an explanation for the oscillations seen in the compressible pressure in the direct numerical simulations of homogeneous shear, and 3) for turbulence closures accounting for the compressibility of velocity fluctuations.

1) Initial conditions consistent with small turbulent Mach number asymptotics are constructed. The importance of consistent initial conditions in the direct numerical simulation of compressible turbulence is dramatically illustrated: spurious oscillations associated with inconsistent initial conditions are avoided, and the fluctuating dilatational field is some two orders of magnitude smaller for a compressible isotropic turbulence. For the isotropic decay it is shown that the choice of initial conditions can change the scaling law for the compressible dissipation.

2) A two-time expansion of the Navier-Stokes equations is used to distinguish compressible acoustic and compressible advective modes. A simple conceptual model for weakly compressible turbulence — a forced linear oscillator — is described. It is shown that the evolution equations for the compressible portions of turbulence can be understood as a forced wave equation with refraction. Acoustic modes of the flow can be amplified by refraction and are able to manifest themselves in large fluctuations of the compressible pressure.

3) The consequences of a small turbulent Mach number expansion for the closure of two covariances appearing in the kinetic energy equation, the pressure-dilatation and the dilatational dissipation, are investigated. Comparisons with different models and a discussion of the results in the context of the homogeneous shear is given. In agreement with observations of DNS of compressible turbulence, the dilatational covariances can not account for the large reductions in the growth of compressible shear layers.

Hu, Changqing, and Chi-Wang Shu: *A discontinuous Galerkin finite element method for Hamilton-Jacobi equations*. ICASE Report No. 98-2, (NASA CR-1998-206903), January 23, 1998, 29 pages. Submitted to Journal of Computational Physics.

In this paper, we present a discontinuous Galerkin finite element method for solving the nonlinear Hamilton-Jacobi equations. This method is based on the Runge-Kutta discontinuous Galerkin finite element method for solving conservation laws. The method has the flexibility of treating complicated geometry by using arbitrary triangulation, can achieve high order accuracy with a local, compact stencil, and are suited for efficient parallel implementation. One and two dimensional numerical examples are given to illustrate the capability of the method.

Lewis, Robert Michael, and Virginia Torczon: *Pattern search methods for linearly constrained minimization*. ICASE Report No. 98-3, (NASA CR-1998-206904), January 23, 1998, 25 pages. Submitted to SIAM Journal on Optimization.

We extend pattern search methods to linearly constrained minimization. We develop a general class of feasible point pattern search algorithms and prove global convergence to a Karush-Kuhn-Tucker point. As in the case of unconstrained minimization, pattern search methods for linearly constrained problems accomplish this without explicit recourse to the gradient or the directional derivative. Key to the analysis of the algorithms is the way in which the local search patterns conform to the geometry of the boundary of the feasible region.

Smith, Ralph C.: *A nonlinear physics-based optimal control method for magnetostrictive actuators*. ICASE Report No. 98-4, (NASA CR-1998-206905), January 23, 1998, 32 pages. Submitted to Journal of Intelligent Material Systems and Structures.

This paper addresses the development of a nonlinear optimal control methodology for magnetostrictive actuators. At moderate to high drive levels, the output from these actuators is highly nonlinear and contains significant magnetic and magnetomechanical hysteresis. These dynamics must be accommodated by models and control laws to utilize the full capabilities of the actuators. A characterization based upon ferromagnetic mean field theory provides a model which accurately quantifies both transient and steady state actuator dynamics under a variety of operating conditions. The control method consists of a linear perturbation feedback law used in combination with an optimal open loop nonlinear control. The nonlinear control incorporates the hysteresis and nonlinearities inherent to the transducer and can be computed offline. The feedback control is constructed through linearization of the perturbed system about the optimal system and is efficient for online implementation. As demonstrated through numerical examples, the combined hybrid control is robust and can be readily implemented in linear PDE-based structural models.

Ristorcelli, J.R.: *A kinematically consistent two-point correlation function*. ICASE Report No. 98-5, (NASA CR-1998-206909), January 30, 1998, 14 pages. Submitted to Physics of Fluids.

A simple kinematically consistent expression for the longitudinal two-point correlation function related to both the integral length scale and the Taylor microscale is obtained. On the inner scale, in a region of width inversely proportional to the turbulent Reynolds number, the function has the appropriate curvature at the origin. The expression for two-point correlation is related to the nonlinear cascade rate, or dissipation

ε , a quantity that is carried as part of a typical single-point turbulence closure simulation. Constructing an expression for the two-point correlation whose curvature at the origin is the Taylor microscale incorporates one of the fundamental quantities characterizing turbulence, ε , into a model for the two-point correlation function. The integral of the function also gives, as is required, an outer integral length scale of the turbulence independent of viscosity. The proposed expression is obtained by kinematic arguments; the intention is to produce a practically applicable expression – in terms of simple elementary functions – that allow an analytical evaluation, by asymptotic methods, of diverse functionals relevant to single-point turbulence closures. Using the expression devised an example of the asymptotic method by which functionals of the two-point correlation can be evaluated is given.

Mavriplis, Dimitri J.: *Multigrid strategies for viscous flow solvers on anisotropic unstructured meshes*. ICASE Report No. 98-6, (NASA CR-1998-206910), January 30, 1998, 27 pages. Submitted to AIAA Journal.

Unstructured multigrid techniques for relieving the stiffness associated with high-Reynolds number viscous flow simulations on extremely stretched grids are investigated. One approach consists of employing a semi-coarsening or directional-coarsening technique, based on the directions of strong coupling within the mesh, in order to construct more optimal coarse grid levels. An alternate approach is developed which employs directional implicit smoothing with regular fully coarsened multigrid levels. The directional implicit smoothing is obtained by constructing implicit lines in the unstructured mesh based on the directions of strong coupling. Both approaches yield large increases in convergence rates over the traditional explicit full-coarsening multigrid algorithm. However, maximum benefits are achieved by combining the two approaches in a coupled manner into a single algorithm. An order of magnitude increase in convergence rate over the traditional explicit full-coarsening algorithm is demonstrated, and convergence rates for high-Reynolds number viscous flows which are independent of the grid aspect ratio are obtained. Further acceleration is provided by incorporating low-Mach-number preconditioning techniques, and a Newton-GMRES strategy which employs the multigrid scheme as a preconditioner. The compounding effects of these various techniques on speed of convergence is documented through several example test cases.

Mavriplis, Dimitri J.: *Directional agglomeration multigrid techniques for high-Reynolds number viscous flows*, ICASE Report No. 98-7, (NASA CR-1998-206911), January 30, 1998, 25 pages. Submitted to AIAA Journal.

The quantification and utilization of coupling effects in a prototypical structural acoustic system are examined in this paper. In typical systems, the coupling mechanisms are manifested in two ways. The first leads to the transfer of energy from an ambient field to an adjacent structure and is often responsible for exogenous structural excitation. The second involves the transfer of energy from the vibrating structure to an adjacent field. This is the source of structure-borne noise and is ultimately the mechanism through which structural actuators are utilized to attenuate noise. The examples presented here demonstrate that in fully coupled systems, both mechanisms should be incorporated to accurately model system dynamics. The examples also illustrate advantages and limitations of compensators which utilize the accurate modeling of the structural coupling.

Rubinstein, Robert, Ye Zhou, and Gordon Erlebacher: *Transport coefficients in rotating weakly compressible turbulence*. ICASE Report No. 98-8, (NASA CR-1998-206916), February 6, 1998, 20 pages. Submitted to Physics of Fluids.

Analytical studies of compressible turbulence have found that compressible velocity fluctuations create both effective fluid transport properties and an effective equation of state. This paper investigates the effects of rotation on compressible turbulence. It is shown that rotation modifies the transport properties of compressible turbulence by replacing the turbulence time scale by a rotational time scale, much as rotation modifies the transport properties of incompressible turbulence. But thermal equilibrium properties are modified in a more complex manner. Two regimes are possible: one dominated by incompressible fluctuations, in which the sound speed is modified as it is in non-rotating turbulence, and a rotation dominated regime in which the sound speed enhancement is rotation dependent. The dimensionless parameter which discriminates between regimes is identified. In general, rotation is found to suppress the effects of compressibility. A novel feature of the present analysis is the use of a non-Kolmogorov steady state as the reference state of turbulence. introduction of such steady states expands the power and utility of analytical turbulence closures to a wider range of problems.

Nowaczyk, Ronald H.: *Perceptions of engineers regarding successful engineering team design*. ICASE Report No. 98-9, (NASA CR-1998-206917), February 6, 1998, 17 pages. To be submitted to Group and Organizational Studies.

The perceptions of engineers and scientists at NASA Langley Research Center toward engineering design teams were evaluated. A sample of 49 engineers and scientists rated 60 team behaviors in terms of their relative importance for team success. They also completed a profile of their own perceptions of their strengths and weaknesses as team members. Behaviors related to team success are discussed in terms of those involving the organizational culture and commitment to the team and those dealing with internal team dynamics. The latter behaviors focused on team issues occurring during the early stages of a team's existence. They included the level and extent of debate and discussion regarding methods for completing the team task and the efficient use of team time to explore and discuss methodologies critical to the problem. The discussion includes a comparison of engineering teams with the prototypical business team portrayed in the literature.

Povitsky, A., and D. Ofengeim: *Numerical study of interaction of a vortical density inhomogeneity with shock and expansion waves*. ICASE Report No. 98-10, (NASA CR-1998-206918), February 6, 1998, 23 pages. To be submitted to Journal of Computational Fluid Dynamics.

We studied the interaction of a vortical density inhomogeneity (VDI) with shock and expansion waves. We call the VDI the region of concentrated vorticity (vortex) with a density different from that of ambience. Non-parallel directions of the density gradient normal to the VDI surface and the pressure gradient across a shock wave results in an additional vorticity. The roll-up of the initial round VDI towards a non-symmetrical shape is studied numerically.

Numerical modeling of this interaction is performed by a 2-D Euler code. The use of an adaptive unstructured numerical grid makes it possible to obtain high accuracy and capture regions of induced vorticity with a moderate overall number of mesh points. For the validation of the code, the computational results are compared with available experimental results and good agreement is obtained.

The interaction of the VDI with a propagating shock wave is studied for a range of initial and induced circulations and obtained flow patterns are presented. The splitting of the VDI develops into the formation of a non-symmetrical vortex pair and not in a set of vortices.

A method for the analytical computation of an overall induced circulation Γ_1 as a result of the interaction of a moving VDI with a number of waves is proposed. Simplified, approximated, expressions for Γ_1 are derived and their accuracy is discussed.

The splitting of the VDI passing through the Prandtl-Meyer expansion wave is studied numerically. The obtained VDI patterns are compared to those for the interaction of the VDI with a propagating shock wave for the same values of initial and induced circulations. These patterns have similar shapes for corresponding time moments.

Chen, Ya-Chin, and Jer-Nan Juang: *A novel approach for adaptive signal processing*. ICASE Report No. 98-11, (NASA CR-1998-206920), February 6, 1998, 29 pages.

Adaptive linear predictors have been used extensively in practice in a wide variety of forms. In the main, their theoretical development is based upon the assumption of stationarity of the signals involved, particularly with respect to the second order statistics. On this basis, the well-known normal equations can be formulated. If high-order statistical stationarity is assumed, then the equivalent normal equations involve high-order signal moments. In either case, the cross moments (second or higher) are needed. This renders the adaptive prediction procedure non-blind. A novel procedure for blind adaptive prediction has been proposed and considerable implementation has been made in our contributions in the past year. The approach is based upon a suitable interpretation of blind equalization methods that satisfy the constant modulus property and offers significant deviations from the standard prediction methods. These blind adaptive algorithms are derived by formulating Lagrange equivalents from mechanisms of constrained optimization. In this report, other new update algorithms are derived from the fundamental concepts of advanced system identification to carry out the proposed blind adaptive prediction. The results of the work can be extended to a number of control-related problems, such as disturbance identification. The basic principles are outlined in this report and differences from other existing methods are discussed. The applications implemented are speech processing, such as coding and synthesis. Simulations are included to verify the novel modelling method.

Carpenter, Mark H., Jan Nordstrom, and David Gottlieb: *A stable and conservative interface treatment of arbitrary spatial accuracy*. ICASE Report No. 98-12, (NASA CR-1998-206921), February 6, 1998, 24 pages. Submitted to Journal of Computational Physics.

Stable and accurate interface conditions are derived for the linear advection-diffusion equation. The conditions are functionally independent of the spatial order of accuracy and rely only on the form of the discrete operator. We focus on high-order finite-difference operators that satisfy the summation-by-parts (SBP) property. We prove that stability is a natural consequence of the SBP operators used in conjunction with the new boundary conditions. In addition, we show that the interface treatments are conservative.

New finite-difference operators of spatial accuracy up to sixth order are constructed: these operators satisfy the SBP property. Finite-difference operators are shown to admit design accuracy (p^{th} -order global accuracy) when $(p - 1)^{th}$ -order stencil closures are used near the boundaries if the physical boundary conditions are implemented to at least p^{th} -order accuracy. Stability and accuracy are demonstrated on the nonlinear Burgers' equation for an twelve-subdomain problem with randomly distributed interfaces.

Wilson, Robert V., Ayodeji O. Demuren, and Mark Carpenter: *Higher-order compact schemes for numerical simulation of incompressible flows*. ICASE Report No. 98-13, (NASA CR-1998-206922), February 13, 1998, 46 pages. Submitted to Journal of Computational Physics.

A higher order accurate numerical procedure has been developed for solving incompressible Navier-Stokes equations for 2D or 3D fluid flow problems. It is based on low-storage Runge-Kutta schemes for temporal discretization and fourth and sixth order compact finite-difference schemes for spatial discretization. The particular difficulty of satisfying the divergence-free velocity field required in incompressible fluid flow is resolved by solving a Poisson equation for pressure. It is demonstrated that for consistent global accuracy, it is necessary to employ the same order of accuracy in the discretization of the Poisson equation. Special care is also required to achieve the formal temporal accuracy of the Runge-Kutta schemes. The accuracy of the present procedure is demonstrated by application to several pertinent benchmark problems.

Arian, Eyal, and Veer N. Vatsa: *A preconditioning method for shape optimization governed by the Euler equations*. ICASE Report No. 98-14, (NASA CR-1998-206926), February 20, 1998, 18 pages. Submitted to TCFD.

We consider a classical aerodynamic shape optimization problem subject to the compressible Euler flow equations.

The gradient of the cost functional with respect to the shape variables is derived with the adjoint method at the continuous level.

The Hessian (second order derivative of the cost functional with respect to the shape variables) is approximated also at the continuous level, as first introduced by Arian and Ta'asan (1996).

The approximation of the Hessian is used to approximate the Newton step which is essential to accelerate the numerical solution of the optimization problem.

The design space is discretized in the maximum dimension, i.e., the location of each point on the intersection of the computational mesh with the airfoil is taken to be an independent design variable.

We give numerical examples for 86 design variables in two different flow speeds and achieve an order of magnitude reduction in the cost functional at a computational effort of a full solution of the analysis partial differential equation (PDE).

Banks, H.T., and P.R. Emeric: *Detection of non-symmetrical damage in smart plate-like structures*. ICASE Report No. 98-15, (NASA CR-1998-206931), February 23, 1998. 23 pages. Submitted to Journal of Intelligent Material Systems and Structures.

A two-dimensional model for in-plane vibrations of a cantilever plate with a non-symmetrical damage is used in the context of defect identification in materials with piezoelectric ceramic patches bonded to their surface. These patches can act both as actuators and sensors in a self-analyzing fashion, which is a characteristic of smart materials. A Galerkin method is used to approximate the dynamic response of these structures. The natural frequency shifts due to the damage are estimated numerically and compared to experimental data obtained from tests on cantilever aluminum plate-like structures damaged at different locations with defects of different depths. The damage location and extent are determined by an enhanced least square identification method. Efficacy of the frequency shift based algorithms is demonstrated using experimental data.

Simpson, Timothy W.: *Comparison of response surface and kriging models in the multidisciplinary design of an aerospace nozzle*. ICASE Report No. 98-16, (NASA CR-1998-206935), February 26, 1998, 25 pages. Submitted to the 7th AIAA/NASA/USAF/ISSMO Symposium of Multidisciplinary Design and Analysis.

The use of response surface models and kriging models are compared for approximating non-random, deterministic computer analyses. After discussing the traditional response surface approach for constructing polynomial models for approximation, kriging is presented as an alternative statistical-based approximation method for the design and analysis of computer experiments. Both approximation methods are applied to the multidisciplinary design and analysis of an aerospike nozzle which consists of a computational fluid dynamics model and a finite element analysis model. Error analysis of the response surface and kriging models is performed along with a graphical comparison of the approximations. Four optimization problems are formulated and solved using both approximation models. While neither approximation technique consistently outperforms the other in this example, the kriging models—using only a constant for the underlying global model and a Gaussian correlation function—perform as well as the second order polynomial response surface models.

ICASE COLLOQUIA
October 1, 1997 – March 31, 1998

Name/Affiliation/Title	Date
Hubona, Geoffrey S., Virginia Commonwealth University “The Effects of Stereoscopic Viewing, Motion, and Object Shadows on Three-Dimensional Visualization”	October 3
Interrante, Victoria, ICASE “Strategies for Effectively Visualizing 3D Flow with Volume Line Integral Convolution”	October 17
Loncaric, Josip, ICASE “Sensor/Actuator Placement Via Optimal Distributed Control of Exterior Stokes Flow”	October 24
Povitsky, Alexander, ICASE “Behavior of Vortical Density Inhomogeneity after Interaction with Shock and Expansion Waves”	November 12
Chen, Po-Shu, University of Colorado at Boulder “A Unified Framework of FETI Solver and Its Applications”	November 17
DeSouza, Fenella, National Research Council, Canada “On the Large Scale Organization of a Turbulent Boundary Layer Disturbed by a Circular Cylinder”	December 15
Karpel, Moti, Technion Israel Institute of Technology “Structural Optimization of Flight Vehicles with Non-Linear Aerodynamic Loads”	January 23
Darmofal, David, Texas A & M University “Robust, Textbook Multigrid Performance for the 2-D Euler Equations Using Semi-Coarsening and Local Partitioning”	January 29
Nocedal, Jorge, Northwestern University “Algorithms and Software for Large-Scale Nonlinear Optimization”	February 6
Abarbanel, Saul, Tel-Aviv University, Israel “Bounded Error Schemes for the Wave Equation on Complex Domains”	February 19

Name/Affiliation/Title	Date
Abarbanel, Saul, Tel-Aviv University, Israel "Perfectly Matched Absorbing Layers for Advective Acoustics"	February 25
Luetttgen, Gerald, Universitaet Passau, Germany "Approaches to State-Space Reduction"	March 30

ICASE STAFF

I. ADMINISTRATIVE

Manuel D. Salas, Director, M.S., Aeronautics and Astronautics, Polytechnic Institute of Brooklyn, 1970.
Fluid Mechanics and Numerical Analysis.

Linda T. Johnson, Office and Financial Administrator

Etta M. Blair, Accounting Supervisor

Barbara A. Cardasis, Administrative Secretary

Shelly M. Johnson, Executive Secretary/Visitor Coordinator

Rachel A. Lomas, Payroll and Accounting Clerk

Emily N. Todd, Conference Manager

Gwendolyn W. Wesson, Contract Accounting Clerk

Leon M. Clancy, Senior System Manager

Bryan K. Hess, Assistant System Manager

Gregory P. Wheeler, System Operator to December 1997

II. SCIENCE COUNCIL

Francine Berman, Professor, Department of Computer Science & Engineering, University of California-San Diego.

Joseph E. Flaherty, Amos Eaton Professor, Departments of Computer Science and Mathematical Sciences, Rensselaer Polytechnic Institute.

Geoffrey Fox, Director, Northeast Parallel Architectural Center, Syracuse University.

David Gottlieb, Professor, Division of Applied Mathematics, Brown University.

Foresster Johnson, Aerodynamics Research, Boeing Commercial Airplane Group.

Robert W. MacCormack, Professor, Department of Aeronautics and Astronautics, Stanford University.

Steven A. Orszag, Professor, Program in Applied & Computational Mathematics, Princeton University.

Stanley G. Rubin, Professor, Department of Aerospace Engineering and Engineering Mechanics, University of Cincinnati.

Manuel D. Salas, Director, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center.

III. RESEARCH FELLOWS

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Applied & Numerical Mathematics [Grid Techniques for Computational Fluid Dynamics]. (February 1997 to August 2001)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Computer Science [Programming Languages for Multiprocessor Systems]. (January 1991 to September 1999)

IV. SENIOR STAFF SCIENTISTS

Thomas W. Crockett - B.S., Mathematics, The College of William & Mary, 1977. Computer Science [System Software for Parallel Computing, Computer Graphics, and Scientific Visualization]. (February 1987 to August 2000)

Sharath S. Girimaji - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1990. Fluid Mechanics [Turbulence and Combustion]. (July 1993 to August 1999)

R. Michael Lewis - Ph.D., Mathematical Sciences, Rice University, 1989. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (May 1995 to August 2000)

Josip Lončarić - Ph.D., Applied Mathematics, Harvard University, 1985. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (March 1996 to August 1999)

J. Ray Ristorcelli - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1991. Fluid Mechanics [Turbulence Modeling]. (December 1996 to August 1999)

David Sidilkover - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1989. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]. (November 1994 to August 1998)

Ye Zhou - Ph.D., Physics, College of William and Mary, 1987. Fluid Mechanics [Turbulence Modeling]. (October 1992 to September 1998)

V. SCIENTIFIC STAFF

Brian G. Allan - Ph.D., Mechanical Engineering, University of California at Berkeley, 1996. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (February 1996 to August 1999)

Eyal Arian - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1995. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (October 1994 to August 1999)

Po-Shu Chen - Ph.D., Aerospace Engineering, University of Colorado-Boulder, 1997. Physical Sciences [Computational Structures]. (January 1998 to January 2001)

Stephen Guattery - Ph.D., Computer Science, Carnegie Mellon University, 1995. Computer Science [Parallel Numerical Algorithms, including Partitioning and Mapping]. (September 1995 to August 1998)

Victoria L. Interrante - Ph.D., Computer Science, University of North Carolina at Chapel Hill, 1996. Computer Science [Scientific Visualization]. (March 1996 to August 1998)

Leland M. Jameson - Ph.D., Applied Mathematics, Brown University, 1993. Applied & Numerical Mathematics [Multiresolution Schemes]. (August 1996 to October 1997)

Li-Shi Luo - Ph.D., Physics, Georgia Institute of Technology, 1993. Computer Science [Parallel Algorithms]. (November 1996 to October 1999)

Kwan-Liu Ma - Ph.D., Computer Science, University of Utah, 1993. Computer Science [Visualization]. (May 1993 to August 1999)

Alexander Povitsky - Ph.D., Mechanical Engineering, Moscow Institute of Steel and Alloys Technology (MISA), Russia, 1988. Computer Science [Parallelization and Formulation of Higher Order Schemes for Aeroacoustics Noise Propagation]. (October 1997 to August 2000)

VI. VISITING SCIENTISTS

Sang-Hyon Chu - M.S., Chemical Engineering, Seoul National University, 1994. Physical Sciences [Smart Materials and Flow Control]. (March 1998 to November 1998)

Remi Draï - M.S., Mathematics, Université de Nice-Sophia Antipolis, France, 1994. Teaching Assistant, Institut du Non linéaire de Nice and Department of Mathematics, Université de Nice-Sophia Antipolis, France. Physical Sciences [Modeling and Control Dealing with the Interaction of Fuel SLOSH with Rigid Body Motion]. (March 1998 to October 1998)

John Van Rosendale - Ph.D., Computer Science, University of Illinois, 1980. Program Director for New Technologies, Division of Advanced Scientific Computing, National Science Foundation. Computer Science [Parallel Computing]. (July 1994 to March 1999)

Hyeon-Ju Yoon - Ph.D., Computer Science, Korea Advanced Institute of Science and Technology, 1997. Postdoctoral Researcher, Department of Computer Science, Korea Advanced Institute of Science and Technology. Computer Science [Computer Graphics and Parallel and Distributed Processing]. (November 1997 to November 1998)

VII. SHORT TERM VISITING SCIENTISTS

Kenneth Eure - M.S., Electrical Engineering, Virginia Polytechnic Institute and State University, 1991. Physical Sciences [Noise Control]. (October 1997 to December 1997)

Maurice Holt - Ph.D., Mathematics, University of Manchester, England, 1948. Professor, College of Engineering, University of California-Berkeley. Applied & Numerical Mathematics. (October 1997)

Ignacio M. Llorente - Ph.D., Computer Science, Complutense University of Madrid, Spain, 1995. Associate Professor, Department of Computer Architecture, Complutense University of Madrid, Spain. Computer Science. (September 1997 to December 1997)

VIII. ASSOCIATE RESEARCH FELLOW

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Computer Science [Parallel Numerical Algorithms]

IX. CONSULTANTS

Saul Abarbanel - Ph.D., Theoretical Aerodynamics, Massachusetts Institute of Technology, 1959. Professor, Department of Applied Mathematics, Tel-Aviv University, Israel. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamics and Aeroacoustic Computations]

Ponnampalam Balakumar - Ph.D., Aeronautics and Astronautics, Massachusetts Institute of Technology, 1986. Associate Professor, Department of Aerospace Engineering, Old Dominion University. Fluid Mechanics [Stability and Transition]

H. Thomas Banks - Ph.D., Applied Mathematics, Purdue University, 1967. Professor, Department of Mathematics, Center for Research in Scientific Computations, North Carolina State University. Applied & Numerical Mathematics [Control Theory]

Richard W. Barnwell - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1968. Professor, Department of Aerospace and Ocean Engineering, Engineering Science and Mechanics. Virginia Polytechnic Institute and State University. Fluid Mechanics [Turbulence Modeling]

Alvin Bayliss - Ph.D., Mathematics, New York University, 1975. Associate Professor, Technological Institute, Northwestern University. Fluid Mechanics [Numerical Solution of the Equations of Fluid Flow and Acoustics]

Oktay Baysal - Ph.D., Mechanical Engineering, Louisiana State University, 1982. Eminent Scholar and Professor, Department of Aerospace Engineering, Old Dominion University. Applied & Numerical Mathematics

Achi Brandt - Ph.D., Mathematics, Weizmann Institute of Science, 1965. Professor, Department of Applied Mathematics, Weizmann Institute of Science, Israel. Applied & Numerical Mathematics [Convergence Acceleration]

Barbara M. Chapman - M.S., Mathematics, University of Canterbury, Christchurch, New Zealand, 1985. Director, European Institute for Parallel Computing, University of Vienna. Computer Science [Parallel Language Extensions and Optimizations for Parallel Compilers]

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineering and Mechanics, Old Dominion University. Fluid Mechanics [Numerical Modeling of Turbulent Flows]

Geoffrey Fox - Ph.D., Physics, Cambridge University, 1967. Professor, Department of Computer Science, Syracuse University. Computer Science [Networking]

David Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Ford Foundation Professor & Chair, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Boundary Conditions for Hyperbolic Systems]

Chester E. Grosch - Ph.D., Physics and Fluid Dynamics, Stevens Institute of Technology, 1967. Professor, Department of Computer Science and Slover Professor, Department of Oceanography, Old Dominion University. Fluid Mechanics [Turbulence and Acoustics]

Jan S. Hesthaven - Ph.D., Applied Mathematics/Numerical Analysis, Technical University of Denmark, 1995. Visiting Assistant Professor, Division of Applied Mathematics, Brown University. Physical Sciences [Computational Electromagnetics]

Fang Q. Hu - Ph.D., Applied Mathematics, Florida State University, 1990. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Aeroacoustics]

Frank Kozusko - Ph.D., Computational and Applied Mathematics, Old Dominion University, 1995. Assistant Professor, Department of Mathematics, Hampton University. Fluid Mechanics [Airfoil Design]

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Asymptotic and Numerical Methods for Computational Fluid Dynamics]

Kurt Maly - Ph.D., Computer Science, Courant Institute, New York University, 1973. Kaufman Professor and Chair, Department of Computer Science, Old Dominion University. Computer Science [High Performance Communication]

James E. Martin - Ph.D., Applied Mathematics, Brown University, 1991. Assistant Professor, Department of Mathematics, Christopher Newport University. Fluid Mechanics [Turbulence and Computation]

Jan Nordstrom - Ph.D., Numerical Analysis, Uppsala University, Sweden, 1993. Senior Scientist, The Aeronautical Research Institute of Sweden. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Alex Pothén - Ph.D., Applied Mathematics, Cornell University, 1984. Professor, Department of Computer Science, Old Dominion University. Computer Science [Parallel Numerical Algorithms]

Robert Rubinstein - Ph.D., Mathematics, Massachusetts Institute of Technology, 1972. Fluid Mechanics [Drag Reduction and Shock Wave Propagation in a Gas Far-From Equilibrium]

Viktor Ryaben'kii - Ph.D., Stability of Difference Equations, Moscow State University, 1953. Leading Research Scientist, Keldysh Institute for Applied Mathematics, Russian Academy of Sciences and Full Professor, Department of Control and Applied Mathematics, Moscow Institute of Physics and Technology. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Chi-Wang Shu - Ph.D., Mathematics, University of California-Los Angeles, 1986. Associate Professor, Division of Applied Mathematics, Brown University. Fluid Mechanics [Computational Aeroacoustics]

Ralph C. Smith - Ph.D., Mathematics, Montana State University, 1990. Assistant Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Optimal Control Techniques for Structural Acoustics Problems]

Siva Thangam - Ph.D., Mechanical Engineering, Rutgers University, 1980. Professor, Department of Mechanical Engineering, Stevens Institute of Technology. Fluid Mechanics [Turbulence Modeling and Simulation]

Christine A. Toomer - Ph.D., Mathematics, Imperial College of Science, Technology & Medicine, London University, 1990. Research Scientist, Aerodynamics & Vulnerability Department, Sowerby Research Centre, British Aerospace, United Kingdom. Applied & Numerical Mathematics [Methods for Aerodynamic Design and Optimization]

Virginia Torczon - Ph.D., Mathematical Sciences, Rice University, 1989. Assistant Professor, Department of Computer Science, The College of William & Mary. Computer Science [Parallel Algorithms for Optimization including Multidisciplinary Design Optimization]

Semyon V. Tsynkov - Ph.D., Computational Mathematics, Keldysh Institute for Applied Mathematics, Russian Academy of Sciences, 1991. Senior Lecturer, Department of Applied Mathematics, Tel-Aviv University, Israel. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

George M. Vahala - Ph.D., Physics, University of Iowa, 1972. Professor, Department of Physics, The College of William & Mary. Fluid Mechanics [Group Renormalization Methods for Turbulence Approximation]

Robert G. Voigt - Ph.D., Mathematics, University of Maryland, 1969. Professor, Computational Science Program, The College of William & Mary. Computer Science [High Performance Computing]

Hans Zima - Ph.D., Mathematics, University of Vienna, Austria, 1964. Professor, Institute for Software Technology and Parallel Systems, University of Vienna, Austria. Computer Science [Compiler Development for Parallel and Distributed Multiprocessors]

Mohammad Zubair - Ph.D., Computer Science, Indian Institute of Technology, Delhi, India, 1987. Professor, Department of Computer Science, Old Dominion University. Computer Science [Performance of Unstructured Flow-Solvers on Multi Processor Machines]

X. GRADUATE STUDENTS

Ahmed H. Al-Theneyan - Department of Computer Science, Old Dominion University. (February 1998 to Present)

Abdelkader Baggag - Department of Computer Science, Purdue University. (September 1995 to Present)

Ya-Chin Chen - Department of Electrical and Electronic Engineering, Imperial College of Science, Technology, and Medicine, London, United Kingdom. (October 1997 to September 1998)

Zhikai Chen - Department of Computer Science, Old Dominion University. (February 1997 to January 1998)

David C. Cronk - Department of Computer Science, The College of William & Mary. (August 1993 to Present)

David Hysom - Department of Computer Science, Old Dominion University. (October 1997 to February 1998)

Nilan Karunaratne - Department of Computer Science, Old Dominion University. (August 1995 to Present)

Dinesh Kaushik - Department of Computer Science, Old Dominion University. (May 1997 to Present)

Gary Kumfert - Department of Computer Science, Old Dominion University. (January 1997 to Present)

Georgia Liu - Department of Computer Science, Old Dominion University. (February 1997 to Present)

Seth D. Milder - Department of Physics and Astronomy, George Mason University. (September 1997 to Present)

Deborah F. Pilkey - Department of Engineering Sciences & Mechanics, Virginia Polytechnic Institute and State University. (October 1995 to Present)

Kevin Roc - Department of Computer Science, The College of William & Mary. (May 1995 to Present)

Ming-Yun Shih - Department of Computer Science, The College of William & Mary. (May 1997 to Present)

Diann P. Smith - Department of Computer Science, The College of William & Mary. (August 1997 to Present)

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY(Leave blank)	2. REPORT DATE July 1998	3. REPORT TYPE AND DATES COVERED Contractor Report		
4. TITLE AND SUBTITLE Semiannual Report. October 1, 1997 through March 31, 1998		5. FUNDING NUMBERS C NAS1-97046 WU 505-90-52-01		
6. AUTHOR(S)				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Computer Applications in Science and Engineering Mail Stop 403, NASA Langley Research Center Hampton, VA 23681-2199		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23681-2199		10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA/CR-1998-208450		
11. SUPPLEMENTARY NOTES Langley Technical Monitor: Dennis M. Bushnell Final Report				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified -Unlimited Subject Category 59 Distribution: Nonstandard Availability: NASA-CASI (301)621-0390		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, fluid mechanics, and computer science during the period October 1, 1997 through March 31, 1998.				
14. SUBJECT TERMS applied mathematics; multidisciplinary design optimization; fluid mechanics; turbulence; flow control; acoustics; computer science; system software; systems engineering; parallel algorithms			15. NUMBER OF PAGES 51	16. PRICE CODE A04
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	